ISCR

Institute for Scientific Computing Research

Fiscal Year 2000 Annual Report

http://www.llnl.gov/casc/iscr/

Lawrence Livermore National Laboratory
P.O. Box 808, L-561, Livermore, CA 94551

CASC

Center for Applied Scientific Computing

UNIVERSITY RELATIONS PROGRAM

The University Relations Program (URP) encourages collaborative research between Lawrence Livermore National Laboratory (LLNL) and the University of California campuses. The Institute for Scientific Computing Research (ISCR) actively participates in such collaborative research, and this report details the Fiscal Year 2000 projects jointly served by URP and ISCR. For a full discussion of all URP projects in FY 2000, please request a copy of the URP FY 2000 Annual Report by contacting

Lawrence Livermore National Laboratory
Christina M. Budwine, University Relations Program
P. O. Box 808, L-413
Livermore, CA 94551

UCRL-LR-133866-00

Work performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

Contents

The Mission of the ISCR

The Institute for Scientific Computing Research (ISCR) at Lawrence Livermore National Laboratory is jointly administered by the Center for Applied Scientific Computing (CASC) and the University Relations Program (URP), and this joint relationship expresses its mission. An extensively externally networked ISCR cost-effectively expands the level and scope of national computational science expertise available to the laboratory through CASC. The URP, with its infrastructure for managing five institutes and numerous educational programs at LLNL, assumes much of the logistical burden that is unavoidable in bridging the laboratory's internal computational research environment with that of the academic community.

As large-scale simulations on the parallel platforms of DOE's Accelerated Strategic Computing Initiative become increasingly important to the overall mission of LLNL, the role of the ISCR expands in importance, accordingly.

Relying primarily on non-permanent staffing, the ISCR complements laboratory research in areas of the computer and information sciences that are needed at the frontier of laboratory missions. The ISCR strives to be the "eyes and ears" of the laboratory in the computer and information sciences, in keeping the laboratory aware of and connected to important external advances. It also attempts to be "feet and hands," in carrying those advances into the laboratory and incorporating them into practice. In addition to conducting research, the ISCR provides continuing education opportunities to laboratory personnel, in the form of on-site workshops taught by experts on novel software or hardware technologies.

The ISCR also seeks to influence the research community external to the laboratory to pursue laboratory-related interests and to train the workforce that will be required by the laboratory. Part of the performance of this function is interpreting to the external community appropriate (unclassified) aspects of the laboratory's own contributions to the computer and information sciences—contributions that its unique mission and unique resources give it a unique opportunity and responsibility to make.

Of the three principal means of packaging scientific ideas for transfer—people, papers, and software—experience suggests that the most effective means is people. The programs of the ISCR are therefore people-intensive.

Finally, the ISCR, together with CASC, confers an organizational identity on the burgeoning computer and information sciences research activity at LLNL and serves as a point of contact within the laboratory for computer and information scientists from outside.

Institute for Scientific Computing Research
Fiscal Year 2000 Director's Report

Large-scale scientific computation, and all of the disciplines that support it and help to validate it, have been placed at the focus of Lawrence Livermore National Laboratory by the Accelerated Strategic Computing Initiative. The Laboratory operates the computer with the highest peak performance in the world and has undertaken some of the largest and most compute-intensive simulations ever performed. However, computers at architectural extremes are notoriously difficult to use efficiently, and successes (such as the Laboratory's two Bell Prizes awarded in November 1999) only point out the need for much better ways of interacting with the results of large-scale simulations.

Advances in scientific computing research have therefore never been more vital to the core missions of the Laboratory than at present. Computational science is evolving so rapidly along every one of its research fronts that to remain on the leading edge, the Laboratory must engage researchers at many academic centers of excellence. In FY 2000, the Institute for Scientific Computing Research (ISCR) has expanded the Laboratory's bridge to the academic community in the form of collaborative subcontracts, visiting faculty, student internships, workshops, and an active seminar series.

ISCR research participants are integrated almost seamlessly with the Laboratory's Center for Applied Scientific Computing (CASC), which, in turn, addresses computational challenges arising throughout the Laboratory. Administratively, the ISCR flourishes under the Laboratory's University Relations Program (URP). Together with the other four Institutes of the URP, it navigates a course that allows the Laboratory to benefit from academic exchanges while preserving national security. While FY 2000 brought more than its share of challenges to the operation of an academic-like research enterprise within the context of a national security laboratory, the results declare the challenges well met and worth the continued effort.

Fiscal year 2000 was the first full year under Acting Director David Keyes. Keyes, the Richard F. Barry Chair of Mathematics & Statistics at Old Dominion University and an ISCR faculty participant since October 1997, dedicated half of his time to the technical program of the ISCR. Jill Dunaway continued as the full-time Institute Administrator. Leslie Bills continued her support on the seminar series and Terry Garrigan came aboard in time to help with the very large summer program and program expansions in other areas, as indicated below.

ISCR Monthly Special Events

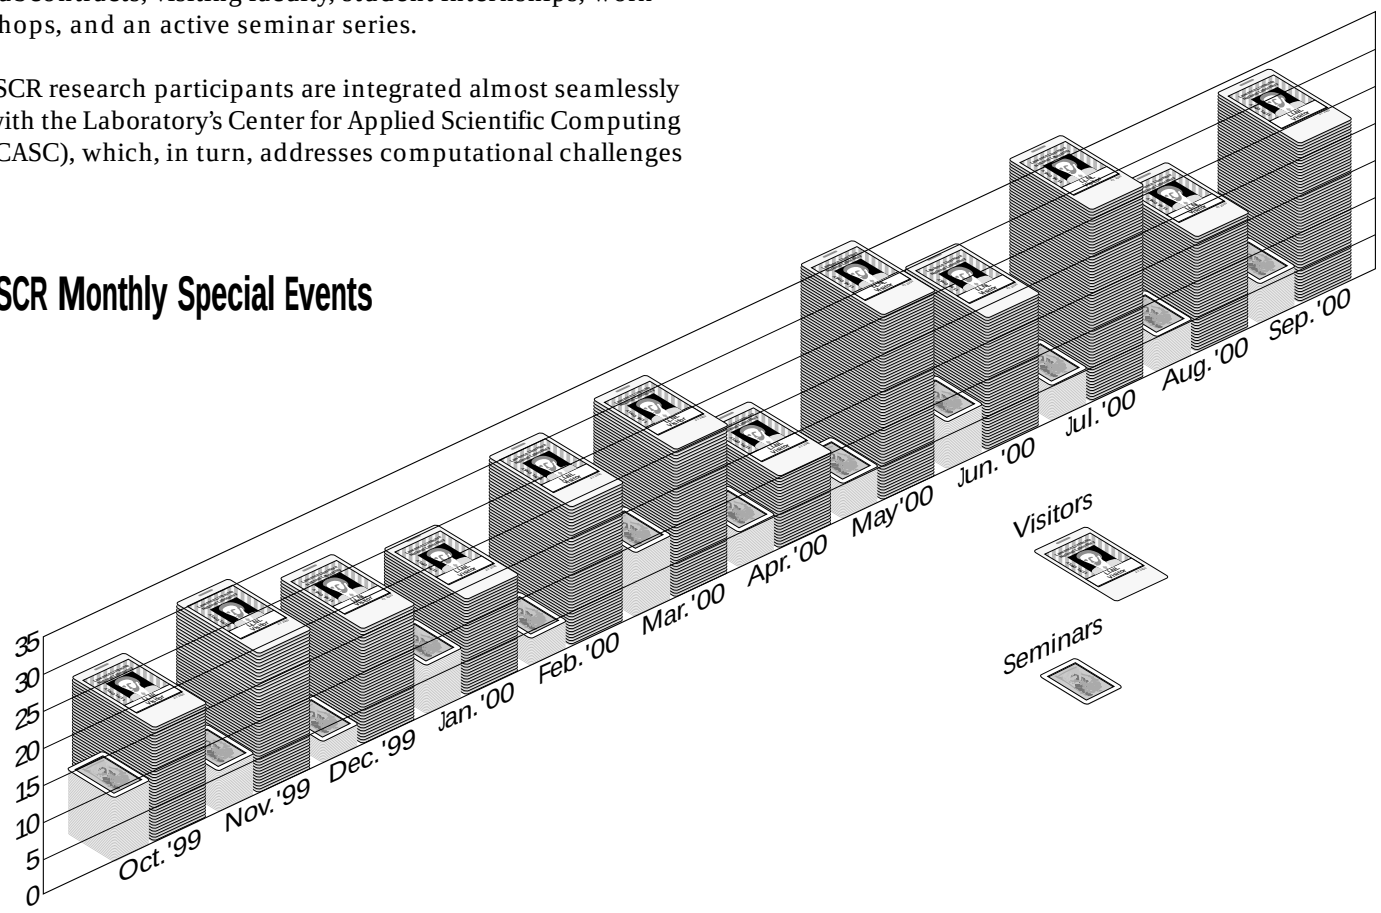

In FY 2000, we launched our ASCI Institute for Terascale Simulation Lecture Series, featuring visits from Fred Brooks, Peter Lax, Burton Smith, and Gilbert Strang. A special section of this annual report is devoted to the abstracts and biosketches of these distinguished lecturers. The ITS Lectures typically draw two to three hundred people from around the laboratory and surrounding scientific community. They are archived on video and available at the LLNL Technical Library. We plan to continue this series with approximately six "movers and shakers" in high-end simulation and its enabling technologies per year.

In April, the ISCR co-sponsored the annual *Copper Mountain Conference*, in Copper Mountain Colorado. The 2000 meeting was devoted to Iterative Methods. Eight members of the CASC scientific staff presented papers, as did twenty of the academic collaborators of the ISCR.

In May, the ISCR organized a three-day *Power Programming Short Course* to enable laboratory code developers (in CASC and in the other divisions) to come to grips, in advance, with the ASCI White machine, of which the laboratory took delivery late in the year. The instructors were Steve White of IBM, Larry Carter of UCSD, David Culler of UCB, Clint Whalley of the University of Tennessee, and Bill Gropp of Argonne. Sixty-five people attended.

In June, with the advent of our large student summer program and sponsorship from the Defense Programs office of DOE HQ, we launched an *Internships in Terascale Simulation Technology* tutorial series, a ten-week series with two-lectures per week. The tutors included three of LLNL's recent computational science textbook authors (Alice Koniges, John May, and Van Henson), five LLNL computational physicists (David Brown, Garry Rodrigue, Howard Scott, Alek Shestakov, and Lin Yang), CASC computer scientist Gary Kumfert, and the ISCR Director. Koniges' *Industrial-Strength Scientific Computing* and May's *Parallel I/O* were published by Morgan-Kaufman during FY 2000. Henson co-authored, with Bill Briggs and Steve McCormick, a recent update of Briggs' "best selling" 1987 monograph, *A Multigrid Tutorial*. Though intended for students, permanent CASC researchers attended an occasional subseries of the lectures.

In July, under the direction of CASC scientist Carol Woodward, the ISCR organized a three-day *Workshop on Solution Methods for Large-scale Nonlinear Problems* at a hotel in nearby Pleasanton. This workshop was capped at 48 attendees, for good discussions. This workshop was an ASCI-context successor to a 1995 workshop with a similar title organized by Professor Homer Walker (now Chair of Mathematical Sciences at WPI). Also dubbed, *The Rootfinders' Ball*, the attendance list of this workshop reads like a "Who's Who" in parallel implicit methods for PDEs. A special issue of Linear Algebra and its Applications with papers drawn from the workshop is being guest-edited by organizer Woodward and CASC colleague Panayot Vassilevski.

We also in July co-sponsored a two-day *Workshop on Mining Scientific Data Sets*. Professor Vipin Kumar, an ASCI Level-2 collaborator, was the host at the Army HPC Research Center at the University of Minnesota, and CASC scientist Chandrika Kamath was a co-organizer. Livermore affiliates provided three of the talks. Most of the principals in the nascent field of scientific data mining were among the 110 in attendance.

Rounding out the very busy month of July, the ISCR co-hosted the *Computational Science Graduate Fellows Conference*, providing local organization and a substantial part of the technical program for this three-day conference. The Krell Institute, which manages the CSGF program for the Department of Energy, made the most of this immersive technical get-together, bringing 38 of their fellows to join the 7 who were already interning at CASC in the ISCR summer student program.

Completing the year's technical meetings, the ISCR co-sponsored the *Fifth Symposium on Overset Grids and Solution Technology* at UC Davis. CASC scientist David Brown was a co-organizer. Six members of CASC and three ISCR affiliates presented work carried out in the laboratory's OVERTURE and SAMRAI frameworks.

In FY 2000, the ISCR brought to the laboratory a vigorous contingent of post-docs, faculty visitors, and students. Twenty faculty visitors were in residence for more than just a seminar visit—for a week to a semester. Eight post-docs made the ISCR their home this past year. We also had 55 students in residence, mostly for 8–10 weeks of the summer, but several of them for a semester or a full year. Each of these students was in a research relationship with one of CASC's 87 full-time technical staff.

Looking ahead, the ISCR anticipates co-sponsorship of post-docs with the new NSF Institute for Pure and Applied Mathematics (IPAM) at UCLA. IPAM co-directors Mark Green, Eitan Tadmor, and Tony Chan consider LLNL's ISCR to be a prime off-campus partner.

The pages of this report summarize the activities of the faculty members, post-doctoral researchers, students, and guests from industry and other laboratories who participated in LLNL's computational mission under the auspices of the ISCR during FY 2000. Altogether, the ISCR hosted 220 visits from 158 different visitors, who gave a total of 72 seminars on site. The vast majority of the visitors were from academia, with 11% from industry and 13% from other laboratories. Visitors from outside of the United States made up 6% of the total. The histograms on page 4 chart the numbers of visitors and seminars as a function of the month of the fiscal year.

Most of the material of this annual report comes directly from the visitors and principal investigators of the projects being reported, who selected formats convenient for their purposes. We thank Dan Moore and Linda Moore of the Technical Information Division of LLNL for their graphic artistry in producing an easily navigated and visually pleasing document.

We hope that you enjoy examining this report on the ISCR's diverse activities in FY 2000. For further information about the Institute, please contact us at the address below. Inquiries about how you might enhance the on-going FY 2001 program at the ISCR, or beyond, are welcome.

David Keyes

David E. Keyes
Acting Director, ISCR
925/ 422-1325
dekeyes@llnl.gov
http://www.llnl.gov/casc/people/keyes/

Jill Dunaway
Administrator, ISCR
925/ 422-7132
dunaway4@llnl.gov
http://www.llnl.gov/casc/people/dunaway/

Institute for Scientific Computing Research

Lawrence Livermore National Laboratory
P. O. Box 808, L-561
Livermore, CA 94551

http://www.llnl.gov/casc/iscr/

ISCR Fiscal Year 2000 in Review

FY 2000 Seminar Series (in reverse chronological order)

Gabriel Wittum, University of Heidelberg ..... September 29, 2000
Klaus Stueben, GMD—Forschungszentrum Informationstechnik GmbH ..... September 28, 2000
Luis Chacon, Los Alamos National Laboratory ..... September 27, 2000
Alfred Inselberg, San Diego Supercomputer Center ..... September 22, 2000
Marty Itzkowitz, Sun Microsystems, Inc. ..... September 15, 2000
Sutanu Sarkar, University of California, San Diego ..... September 13, 2000
Gundolf Haase, Johannes Kepler University, Linz ..... September 11, 2000
Paul Reynolds, University of Virginia ..... August 28, 2000
Ed Seidel, Max-Planck-Institut fuer Graviationsphysik ..... August 25, 2000
Ariel Shamir, University of Texas, Austin ..... August 10, 2000
Gregory Balls, University of California, Berkeley ..... August 9, 2000
Randy Bank, University of California, San Diego ..... August 8, 2000
Michael Holst, University of California, San Diego ..... August 1, 2000
Beth Anne Bennett, Yale University ..... July 25, 2000
Zhiqiang Cai, Purdue University ..... July 24, 2000
Steven Allmaras, Boeing ..... July 19, 2000
Gregory Forest, University of North Carolina ..... June 27, 2000
Saul Abarbanel, Tel Aviv University ..... June 14, 2000
Craig Douglas, University of Kentucky ..... June 12, 2000
Barton Miller, University of Wisconsin ..... June 12, 2000
Sean Peisert, San Diego Supercomputer Center ..... June 8, 2000
Kenneth Powell, University of Michigan ..... June 5, 2000
Stephan Knapek, University of Bonn ..... May 31, 2000
Anthony Skjellum, MPI Software Technologies, Inc. ..... May 25, 2000
Eitan Tadmor, University of California, Los Angeles ..... May 18, 2000
Eitan Tadmor, University of California, Los Angeles ..... May 18, 2000
Iain Duff, CCLRC Rutherford Appleton Laboratory ..... May 9, 2000
Karl Warnick, University of Illinois, Urbana-Champaign ..... April 28, 2000
Gerhard Zumbusch, University of Bonn ..... April 25, 2000
X. Sherry Li, Lawrence Berkeley Laboratory ..... April 13, 2000
Padhriac Smyth, University of California, Irvine ..... April 6, 2000
Omar Ghattas, Carnegie Mellon University ..... March 29, 2000
Jan Hesthaven, Brown University ..... March 28, 2000
Jan Hesthaven, Brown University ..... March 27, 2000
Rob Van Der Wijngaart, MRJ Technology Solutions, NASA Ames Research Center ..... March 24, 2000
Joel Saltz, University of Maryland ..... March 23, 2000
N. Radhakrishnan and Raju Namburu, U.S. Army Research Laboratory (ARL) ..... March 16, 2000
Eric DeSturler, University of Illinois, Urbana-Champaign ..... March 14, 2000
Gil Chita, TakeFive Software ..... March 13, 2000

FY 2000 Seminar Series (in reverse chronological order) continued:

Luiz De Rose, IBM TJ Watson Research Center ..... March 7, 2000
Harvey Wasserman, Los Alamos National Laboratory ..... March 3, 2000
Dimitri Mavriplis, ICASE NASA Langley Research Center ..... March 2, 2000
Inez Heinz, Lawrence Livermore National Laboratory ..... February 28, 2000
Gabriel Wittum, University of Heidelberg ..... February 24, 2000
Gerald Hedstrom, LLNL Retiree ..... February 23, 2000
David Padua, University of Illinois ..... February 11, 2000
Hanan Samet, University of Maryland, College Park ..... January 28, 2000
John Quagliano, Los Alamos National Laboratory ..... January 27, 2000
John Quagliano, Los Alamos National Laboratory ..... January 27, 2000
Allen Malony, University of Oregon ..... January 24, 2000
Steve Chapin, Syracuse University ..... January 19, 2000
Alok Choudhary, Northwestern University ..... January 14, 2000
Petter Bjorstad, University of Bergen, Norway ..... January 13, 2000
Art Mirin, Lawrence Livermore National Laboratory ..... January 11, 2000
Linda Stals, Old Dominion University ..... January 6, 2000
Jong Kim, Pohang University of Science and Technology ..... December 14, 1999
Michael Donaldson, Merant ..... December 13, 1999
Sally McKee, University of Utah ..... December 13, 1999
Garth Gibson, Carnegie Mellon University ..... December 2, 1999
V. Ralph Algazi, University of California, Davis ..... November 23, 1999
Jeremy Siek, University of Notre Dame ..... November 22, 1999
Andrew Lumsdaine, University of Notre Dame ..... November 22, 1999
Peter Vanderbilt, MRJ Technology Solutions ..... November 12, 1999
Jeff Hollingworth, University of Maryland, College Park ..... November 2, 1999
Pat Miller, Scientific Computing Applications Division ..... October 20, 1999
Michael Burl, Jet Propulsion Laboratory ..... October 18, 1999
Roy Hemker, University of California, Los Angeles ..... October 15, 1999
Dean Dauger, University of California, Los Angeles ..... October 7, 1999
David Young, The Boeing Company ..... October 5, 1999
Timothy Kelley, North Carolina State University ..... October 4, 1999

FY 2000 Institute for Terascale Simulation Lecture Series

Frederick P. Brooks, University of North Carolina ..... August 30, 2000
Peter Lax, Courant Institute of the Mathematical Sciences, New York University ..... June 9, 2000
Burton Smith, Tera Computer Company ..... May 12, 2000
Gilbert Strang, Massachusetts Institute of Technology ..... April 12, 2000

Visiting Faculty, Guests, Consultants, and Researchers

Visiting and Collaborating Professors

Xiao-Chuan Cai, University of Colorado
David Dean, Front Range Scientific Computations, Inc.
Jack Dongarra, University of Tennessee
Craig Douglas, University of Kentucky
Alejandro Garcia, San Jose State University
Michael Griebel, University of California, San Diego
Nicholas Hadjiconstantinou, Massachusetts Institute of Technology
Michael Holst, University of California, San Diego
Kenneth Joy, University of California, Davis
Karen Karavanic, Portland State University
David Keyes, Old Dominion University
Raytcho Lazarov, Texas A&M University
Byung Lee, University of Vermont
Dimitri Mavriplis, ICASE, NASA Langley Research Center
James McWilliams, University of California, Los Angeles
Christoph Pflaum, University of Wurzburg
Calvin Ribbens, Virginia Polytechnic Institute and State University
John Ruge, Front Range Scientific Computations, Inc.
Donald Schwendeman, Rensselaer Polytechnic Institute
Padhraic Smyth, University of California, Irvine
Robert Snapp, University of Vermont

Participating Guests

Mark Adams, University of California, Berkeley
Marsha Berger, New York University
William Bosl, Stanford University
Marian Brezina, University of Colorado
George Byrne, Illinois Institute of Technology
Roger Crawfis, Ohio State University
David Dean, University of Colorado
Eric de Sturler, University of Illinois
John Fitzgerald, Lawrence Livermore National Laboratory (retired)
Kyle Gallivan, Florida State University
Michael Gertz, University of California, Davis
Michael Griebel, University of Bonn
Bernd Hamann, University of California, Davis
Ulf Hannebutte, Intel Corporation
Stanley Johnson, Lehigh University

Participating Guests (continued)

Kenneth Joy, University of California, Davis
Johannes Kraus, University of Leoben
Raytcho Lazarov, Texas A&M University
Andrea Malagoli, University of Chicago
Michael Minion, University of North Carolina
Joseph Pasciak, Texas A&M University
Michael Pernice, University of Utah
Elbridge Gerry Puckett, University of California, Davis
John Rice, University of California, Berkeley
Ulrich Ruede, University of Erlangen
Thomas Rutaganira, American River College
Yousef Saad, University of Minnesota
Paul Saylor, University of Illinois
Daniel Schikore, Computational Engineering, International
Rob Van Der Wijngaart, NASA Ames Research Center
Gabriel Wittum, University of Heidelberg
Daniel Wolitzer, California State University, Hayward
Ytha Yu, California State University, Hayward
Ludmil Zikatanov, Pennsylvania State University
Gerhard Zumbusch, University of Bonn

Consultants

Bernie Alder, University of California (Professor Emeritus)
Randolph Bank, University of California, San Diego
Leo Breiman, University of California, Berkeley
Harry Dwyer, University of California, Davis
Anne Greenbaum, University of Washington
Chuck Hansen, University of Utah
David Keyes, Old Dominion University
Heinz-Otto Kreiss, University of California, Los Angeles
Thomas Manteuffel, University of Colorado
Stephen McCormick, University of Colorado
Linda Petzold, University of California, Santa Barbara
Steve Schaffer, New Mexico Tech
Homer Walker, Worcester Polytechnic Institute

Department of Applied Science Faculty

Nelson Max
Garry Rodrigue

Postdoctoral Researchers

Robert Anderson
Erick Cantu-Paz
Petri Fast
Jean-Luc Fattebert
Raymond Fellers
Imola Fodor
Barry Lee
Luc Machiels
Brian Miller
Thomas Rutaganira
Leonid Tsap

University Collaborative Research Program Subcontractors

Scott Baden, University of California, San Diego
Jackson Beatty, University of California, Los Angeles
John Dawson, University of California, Los Angeles
Jeffrey Gregg, University of California, Davis
B. S. Manjunath, University of California, Santa Barbara
Warren Mori, University of California, Los Angeles
Linda Petzold, University of California, Santa Barbara
Joachim Raeder, University of California, Los Angeles
Andrew Szeri, University of California, Berkeley

LDRD Project Investigators

Mark Duchaineau, LLNL, Center for Applied Scientific Computing
Chandrika Kamath, LLNL, Center for Applied Scientific Computing

Students

Student Guests

Nathan Crane, University of Illinois
Matt Giamporcaro, Boston University
Charles Hindman, University of Colorado
Jason Hunt, University of Michigan
McKay Hyde, California Institute of Technology
David Hysom, Old Dominion University
Lars Karlsson, Chalmers University of Technology

Student Guests (continued)

Falko Kuester, University of California, Davis
Diem Phuong Nguyen, University of Utah
Stefan Nilsson, Chalmers Institute of Technology
Christopher Oehmen, University of Tennessee
Wing Yee, University of Utah

Department of Applied Science Students

Paul Covello
Rebecca Darlington
Ana Iontcheva
Joseph Koning
Daniel Laney
Sean Lehman
Tim Pierce
Jonathan Rochez
Subhasis Saha
Bahrad Sokhansanj
Jay Thomas
Michael Wickett

ISCR Students

Marcel Arndt, University of Bonn
Travis Austin, University of Colorado
Zachary Belanger, Oakland University
Martin Bertram, University of California, Davis
Melvina Blackgoat, Northern Arizona University
Kathleen Bonnell, University of California, Davis
Timothy Chartier, University of Colorado
Tom Dossa, Santa Clara University
Jochen Garcke, University of Bonn
Aaron Herrnstein, University of California, Davis
Chisup Kim, Texas A&M University
Imelda Kirby, University of Washington
Stephan Knapek, University of Bonn
Frank Koster, University of Bonn
David Nault, University of Cincinnati

ISCR Students (continued)

Serban Porumbescu, University of California, Davis
Robert Rieben, University of California, Davis
M. Alex Schweitzer, University of Bonn
Danny Thorne, University of Kentucky
Stanimire Tomov, Texas A&M University
Clinton Torres, Northern Arizona University
Serge van Criekingen, Northwestern University
Kevin Vlack, University of Illinois

ITST Students

Lucas Ackerman, Worcester Polytechnic Institute
Brian Ball, Worcester Polytechnic Institute
Janine Bennett, University of California, Davis
Richard Cook, University of California, Davis
Michael Flanagan, Texas A&M University
David Hysom, Old Dominion University
Linh Lieu, University of California, Davis
David Littau, University of Minnesota
Michael McCracken, Penn State University
Jason Morgan, University of Utah
Joshua Senecal, University of California, Davis

National Physical Science Consortium (NPSC) Students

Rachel Karchin, University of California, Santa Cruz
Imelda Kirby, University of Washington
Megan Thomas, University of California, Berkeley

Workshops and Conferences

Common Component Architecture (CCA) Forum, Oakland, CA, March 2000
Copper Mountain Conference, Copper Mountain, CO, April 2000
Power Programming Short Course, Livermore, CA, May 2000
Workshop on Mining Scientific Datasets, Minneapolis, MN, July 2000
Solution Methods for Large-Scale Nonlinear Problems, Pleasanton, CA, July 2000
CSGF Conference, Livermore, CA, July 2000
5th Symposium on Overset Grids & Solution Technology, Davis, CA, September 2000

ISCR

Institute for Scientific Computing Research

Seminar Series Abstracts

(in reverse chronological order)

CASC Center for Applied Scientific Computing

September 29, 2000

Multigrid Methods for Porous Media Flows

Gabriel Wittum
University of Heidelberg

Email: wittum@iwr.uni-heidelberg.de

Abstract:

Porous media flows occur in many situations, e.g., flow of ground and soil water, catalyzers, filters, flow through human tissue and skin, and many more. Modeling single and multi-phase flow through porous media is a challenging mathematical topic on account of their multi-scale character.

The simulation of such problems has to deal with many difficulties, such as heterogenous materials resulting in stochastically jumping coefficients, non-linear problems with singular coefficient functions, complicated geometries, and, in particular for geological problems, strong anisotropies, convection dominated problems, and many more. All these problems are serious difficulties for multigrid methods. The most interesting topic in applying multigrid to porous media flow is how to relate the coarse-grid problems with the modelling scale.

In the lecture we show multigrid strategies for porous media flow problems. We show their succesful application to a number of different applications.

Research web page: http://www.iwr.uni-heidelberg.de/~techsim/

Institution web page: http://www.uni-heidelberg.de/index_e.html

September 28, 2000

Algebraic Multigrid In An Industrial Environment

Klaus Stueben

GMD - Forschungszentrum Informationstechnik GmbH

Email: stueben@gmd.de

Abstract:

Algebraic multigrid (AMG) has been shown to be very efficient and robust for the solution of various types of linear algebraic systems of equations, in particular those arising from the discretization of scalar partial differential equations. However, major research is still required to extend AMG's range of applicability. For instance, its applicability to systems of partial differential equations has not yet reached a robustness and efficiency comparable to that of the scalar case. But even for certain scalar applications, the performance of AMG may substantially deteriorate.

In principle, AMG is a highly interesting candidate for a "plug-in" solver in industrial codes. However, much work has to be invested to turn an academic code into one which satisfies industrial requirements. Rapid convergence is only one of these requirements. Much more important, however, are robustness and low memory requirements. The original AMG (AMG1R5) was far from really meeting these requirements to a satisfactory extent.

We will give a review on some of our experiences which led to the development of a completely new AMG code, SAMG. The underlying algorithm, although still close to the original AMG ideas, contains various enhancements and has been integrated as a solver into several commercial codes (from computational fluid dynamics, oil-reservoir simulation, and process simulation in semi-conductor applications). Developments regarding SAMG's integration into a commercial structural mechanics code are ongoing. We will discuss critical situations and present remedies for some particular scalar and systems applications.

Speaker's web page: http://www.gmd.de/SCAI/people/stueben.html

Institution web page: http://www.gmd.de/Welcome.en.html

September 27, 2000

An Implicit, Nonlinear Newton–Krylov Resistive MHD Solver

Luis Chacon
Los Alamos National Laboratory

Email: chacon@lanl.gov

Abstract:

Implicit time differencing of the resistive magnetohydrodynamic (MHD) equations can step over the limiting time scales—such as Alfven time scales—to resolve the dynamic time scales of interest. However, nonlinearities present in these equations make an implicit implementation cumbersome. Here, viable paths for an implicit, nonlinear time integration of the MHD equations are explored using a 2D reduced viscous-resistive MHD model. The implicit time integration is performed using the Newton–Raphson iterative algorithm, employing Krylov iterative techniques for the required algebraic matrix inversions, implemented Jacobian-free (i.e., without ever forming and storing the Jacobian matrix). Convergence in Krylov techniques is accelerated by preconditioning the initial problem. A "physics based" preconditioner, based on an operator-split approximation to the original set of partial differential equations, is employed. The preconditioner employs low-complexity multigrid techniques to invert approximately the resulting elliptic algebraic systems. The resulting 2D reduced resistive MHD implicit algorithm is shown to be successful in dealing with large time steps (100-250 times the explicit Alfven CFL limit) and fine grids (up to 256×256). The algorithm is second-order accurate in time and efficient. Specifically, the number of Krylov iterations per Newton iteration scales very weakly with the total number of mesh points, and the CPU time scales as the time step to the –0.7 power.

Institution web page: http://www.lanl.gov/

September 22, 2000

Visualizing Multidimensional Geometry with Applications to Multivariate Problems

Alfred Inselberg
San Diego Supercomputer Center

Email: aiisreal@math.tau.ac.il

Abstract:

People working on multivariate (multidimensional) problems will benefit by understanding the underlying geometry, that is, learning what is possible and what is not. For example, in 1917 the physicist Ehrenfest showed that planetary orbits are stable only in dimension 3. Another dimensionality result is that rotating bodies have an axis of rotation only in odd-integer dimensions. The applications presented here will be more down to earth!

With a system of parallel coordinates a one-to-one mapping between subsets of *N*-space and subsets of 2-space is obtained. Lines in *N*-space are represented by *N-1* indexed points. In fact all *p*-flats (planes of dimension *p* in *N*-space) are represented by indexed points where the number of indices is one more than the object's dimension. The representations are generalized to enable the visualization of polytopes and certain kinds of hypersurfaces as well as Convex Objects in *R^N*. Synthetic constructions algorithms involving intersections, proximity, an interior point algorithm and "Line Topologies" of interest in Computer Vision will be presented. There will be interactive demonstrations of Biomedical applications, Process Control, Visual Data Mining (Yields in VLSI production, Finance, Retailing, Feature Extraction from LandSat Data etc.), then Collision Avoidance Algorithms for Air Traffic Control. A visual and computational model of a hypersurface suitable for Optimization, Trade-Off Analysis, and Decision support will be included.

Speaker's web page: http://www.sdsc.edu/~nerona/Inselberg_A2/

Institution web page: http://www.sdsc.edu/

September 15, 2000

Using the Forte Developer Performance Tools

Marty Itzkowitz
Sun Microsystems, Inc.

Email: martyi@eng.sun.com

Abstract:

The Forte Developer 6 performance tools are designed to help answer questions about application performance. In this talk we will discuss the kinds of performance questions that users typically ask, and the user-model for how the tools can help answer them. We will then discuss some specific type of questions for threaded applications, and for Parallel FORTRAN applications, and go through some examples of performance problems that the tools can help address.

For most applications, the basic question users have about performance is "What can I change to improve performance?" That translates into knowing what resources are being used, where in the program they are being used, and how the program got to that point. The design goal of the Forte Developer tools was to minimize the number of mouse-clicks to get the the "A-ha" point.

The user model for the tools is to build the application exactly as one would for production, including any optimization and parallelization. To see annotated source displays, use the -g flag for compilation (-g no longer disables optimization and parallelization). Then dbx, or the debugger GUI, can be used to run the application and measure its performance. From the recorded data the analyzer can present a list of functions, ordered by performance cost, and the callers and callees of any function, with their performance costs attributed to the callers and callees. The analyzer can also show annotated source displays, with per-line performance data, and annotated disassembly displays, with per-instruction data. In addition, the source and disassembly displays can show commentary from the compiler.

For multithreaded executables, users also want to know how the threads are being used, and what costs are involved in synchronizing them. For Parallel FORTRAN applications, the users also want to know which loops are parallel, and which serial, and what the obstacles are to parallelization of serial loops. The users are also interested in the efficiency of load-balancing in parallel regions. The tools provide data to answer all of these questions.

Research web page: http://www.sun.com/forte/fortran/features.html

Institution web page: http://www.sun.com/

September 13, 2000

Direct Simulation of the High-Speed, Turbulent Reacting Shear Layer: Compressibility Effects

Sutanu Sarkar

Mechanical and Aerospace Engineering, UC San Diego

Email: sarkar@mechanics.UCSD.edu

Abstract:

In technological applications such as high-speed propulsion and energetic devices, the variation in thermodynamic variables associated with large heat release and/or high speeds interacts with and modifies the underlying turbulent flow. Direct numerical simulations of the reacting shear layer are performed over a wide range of heat release rates and convective Mach numbers to quantify and understand some of these modifications to the turbulence evolution and structure. Simulations are performed with up to 20 million grid points, high-order discretization, and with large evolution times required to achieve full-blown turbulence. The mixing rates and turbulence structure are found to be substantially different with respect to the incompressible case. DNS results and supporting analysis will be presented to explain the observed compressibility effects.

Speaker's web page: http://www.mae.ucsd.edu/GRAD_BROCH/fluids.htm#SARKAR/

Institution web page: http://www.usd.edu/

September 11, 2000

Special Solution Strategies Inside a Spectral Element Ocean Model

Gundolf Haase
Johannes Kepler University, Linz

Email: ghaase@numa.uni-linz.ac.at

Abstract:

The weather forecast for a few days is rather accurate nowadays, but as soon as the prediction period reaches one month the accuracy of the prediction drops. This is because long-term weather is made in and above the ocean. Collecting more data from the atmosphere above the ocean would help to improve the forecast in coast regions, but long-term weather forecast also requires information on what's going on in the ocean.

I will focus my presentation on a ocean model that substitutes the 3D model by five 2D layers. Therein a filtering of the velocity vector requires the solution of the Laplacian equation in each layer for both velocity components. The discretization uses quadrilateral finite elements (cells) with 7th order test functions. The special difficulty therein consist in the huge amount of cells that restricts the range of applicable numerical methods with respect to memory usage.

I will present several techniques that could be used to solve the resulting system of equations by taking into account specialties of the cells. There are three approaches:

- Schur complement methods with a BPS-like preconditioner.
- Sparse approximation of the original matrix and use of algebraic multigrid methods.
- A Two- or Multigrid method with patch smoothing.

The element preconditioning technique by Stefan Reitzinger can be used in these approaches. Finally, the extension of the approaches to 3D is discussed.

Speaker's web page: http://www.numa.uni-linz.ac.at/Staff/ghaase/haase.html

Institution web page: http://www.uni-linz.ac.at/

August 28, 2000

Cracking Multi-Representation Modeling

Paul Reynolds
University of Virginia

Email: pfr@cs.virginia.edu

Abstract:

Plug-and-Play rarely works for multi-simulations—simulations composed of multiple stand-alone simulations. Among the obvious problems are semantically incompatible simulation components and mismatched units (e.g., of measure). We address a more difficult problem: maintaining a valid correspondence between representations of the same phenomenon in different component simulations when concurrent interactions occur. Correctness of a multi-simulation typically requires that a high degree of consistency be maintained between different representations.

Our approach has been to establish fundamental observations pertaining to the correctness of multi-simulations, and then to establish actions that support the construction of a correct multi-simulation, including the identification of mapping functions and the construction of attribute dependency graphs. We have defined a complete algorithm for constructing a correct multi-simulation. We have tested our approach on a small set of multi-simulations, mainly from the DoD, using each multi-simulation's respective DoD High Level Architecture Federated Object Model as our primary source material. Our approach has performed well in these tests.

We present our approach in depth, acknowledging that we have plucked mostly low-hanging fruit—generalizing and exploiting that which is more easily generalized and exploited. As a result, if you can't deliver mapping functions for maintaining consistency among overlapping representations, we can't necessarily deliver a correct multi-simulation. Are there properties of models that would allow a relaxation of our assumptions? Must every useful solution depend heavily on model semantics? The success of our presentation will be a function of the liveliness of the discussion it engenders on these and related questions.

Speaker's web page: http://www.cs.virginia.edu/~pfr/

Research web page: http://www.cs.virginia.edu/~isotach/

Institution web page: http://www.virginia.edu/

August 25, 2000

Black Holes and Collaborative Computational Science

Ed Seidel

Max-Planck-Institut fuer Gravitationsphysik

Email: eseidel@aei-potsdam.mpg.de

Abstract:

Large-scale simulation problems, such as solving Einstein's equations of general relativity to simulate collisions of black holes, are driving the development of a new generation of advanced computational tools to enable scientists and engineers to work together to exploit "Grid-based" networks of parallel computers to solve problems on scales never before possible. I will discuss recent developments in this area, including the Cactus Computational Toolkit for parallel computing, the Globus Metacomputing Toolkit, and their applications to some important problems in astrophysics and relativity, including large-scale simulations of colliding black holes. Movies of these simulations will be shown. The techniques and toolkits discussed here are designed to address a large class of scientific and engineering applications.

Research web page: http://wwwcactuscode.org/

Institution web page: http://www.aei-potsdam.mpg.de/

August 10, 2000

Multi-Resolution Dynamic Meshes with Arbitrary Deformations

Ariel Shamir
University of Texas, Austin

Email: arik@ticam.utexas.edu

Abstract:

Multi-resolution techniques and models have been shown to be effective for the display and transmission of large static geometric objects. Dynamic environments with internally deforming models and scientific simulations using dynamic meshes pose greater challenges in terms of time and space, and the need the development of similar solutions. In this talk I will introduce the T-DAG, a new adaptive multi-resolution representation for dynamic meshes with arbitrary deformations which include attribute, position, connectivity, and topology changes. T-DAG stands for Time-dependent Directed Acyclic Graph, which defines the structure supporting this representation. I will present an incremental algorithm (in time) for constructing the T-DAG representation of a given input mesh. This enables the traversal and use of the multi-resolution dynamic model for partial playback while still constructing new time-steps. The research is a joint work with V. Pascucci and C.L. Bajaj to appear on IEEE Visualization 2000.

Speaker's web page: http://www.ticam.utexas.edu/~arik/

Institution web page: http://www.utexas.edu/

August 9, 2000

A Finite Difference Domain Decomposition Method Using Local Corrections for the Solution of Poisson's Equation

Gregory Balls
University of California, Berkeley

Email: gballs@cs.berkeley.edu

Abstract:

A domain decomposition method will be presented for computing finite difference solutions to the Poisson equation with infinite domain boundary conditions. This method is a finite difference analogue of Anderson's Method of Local Corrections. The solution is computed in three steps. First, fine grid solutions are computed in parallel using infinite domain boundary conditions on each subdomain. Second, information is transferred globally through a coarse grid representation of the charge, and a global coarse grid solution is found. Third, a fine grid solution is computed on each subdomain using boundary conditions set with the global coarse solution, corrected locally with fine grid information from nearby subdomains.

There are three important features of this algorithm. First, the method requires only a single iteration between the local fine grid solutions and the global coarse representation. Second, the error introduced by the domain decomposition is small relative to the truncation error of a standard single-grid solution. Third, the computed solution is second-order accurate and only weakly dependent on the coarse grid spacing and the number of subdomains. As a result of these features, we are able to compute accurate solutions in parallel with a much smaller ratio of communication to computation than more traditional domain decomposition methods.

Results from implementations of this algorithm on an IBM SP2 and a Cray T3E will be presented. These results verify the overall accuracy and demonstrate the parallel scalability of the method.

Speaker's web page: http://www.cs.berkeley.edu/~gballs/

Institution web page: http://www.berkeley.edu/

August 8, 2000

An Algebraic Multilevel Multigraph Algorithm

Randy Bank
University of California, San Diego

Email: reb@sdna1.ucsd.edu

Abstract:

We describe an algebraic multilevel multigraph algorithm. Many of the multilevel components are generalizations of algorithms originally applied to general sparse Gaussian elimination. Indeed, general sparse Gaussian elimination with minimum degree ordering is a limiting case of our algorithm. Our goal is to develop a procedure which has the robustness and simplicity of use of sparse direct methods, yet offers the opportunity to obtain the optimal or near-optimal complexity typical of classical multigrid methods.

Speaker's web page: http://scicomp.ucsd.edu/~reb/

Institution web page: http://www.ucsd.edu/

August 1, 2000

The Constraints in the Einstein Equations: Well-posedness, a Priori and a Posteriori Error Estimates, and Some Numerical Solutions

Michael Holst
University of California, San Diego

Email: mholst@math.ucsd.edu

Abstract:

In this talk we consider a coupled nonlinear elliptic system representing the Hamiltonian and momentum constraints in the Einstein equations. This system must be solved exactly or numerically to produce consistent initial data for general relativistic simulations of black hole and neutron star collisions. Moreover, the constraints must hold at all times in dynamical situations. Well-posedness of the system on connected compact Riemannian manifolds with Lipschitz boundaries was previously unstudied, and therefore in the first part of the talk we establish that the constraints have unique weak solutions under minimal smoothness assumptions on the data. The proof technique (Riesz-Schauder theory and convex analysis) allows for some degree of negative conformal scalar curvature. We also establish two quasi-optimal a priori error estimates for Galerkin approximations, and we derive an a posteriori error estimate which leads to two distinct error indicators for adaptive simplex subdivision algorithms (one estimator is purely residual based, the other involves a linearized dual problem).

In the second part of the talk we outline an implementation of the adaptivity techniques using the adaptive finite element software package "MC," which is designed to adaptively solve general nonlinear systems of tensor equations on manifolds. MC is an adaptive multilevel finite element package which employs a posteriori error estimation, simplex subdivision, algebraic multilevel methods, global inexact Newton methods, and numerical continuation methods, for the solution of coupled elliptic systems on 2- and 3-manifolds. We demonstrate the capabilities of MC by generating initial data for several different types of black hole spacetimes, including a two-hole collision. To perform these calculations, MC uses a completely new low-communication algorithm for parallel adaptive finite element methods, which was developed jointly with R. Bank at UCSD.

Speaker's web page: http://scicomp.ucsd.edu/~mholst/

Institution web page: http://www.ucsd.edu/

July 25, 2000

Modelling Axisymmetric Laminar Flames: An Application of Local Rectangular Refinement Solution-Adaptive Gridding

Beth Anne Bennett
Yale University

Email: beth.bennett@yale.edu

Abstract:

Over the past two decades, scientific interest in multidimensional laminar and turbulent combustion phenomena has greatly increased, mainly because of tightened government restrictions on pollutant emissions, along with a growing public awareness of environmental concerns. Although today's computers are still not powerful enough to solve exact numerical models of three-dimensional, unsteady, turbulent combustion, a class of problems which *can* be solved is that of chemically complex, multidimensional laminar flames. Because reaction zone thicknesses in such flames can be less than one-hundredth of the overall flame height, the demands on CPU time and memory can be drastically reduced (and the computations made feasible) with the application of a sophisticated adaptive gridding technique.

The results presented here have been obtained using the local rectangular refinement (LRR) solution-adaptive gridding method, which automatically produces orthogonal unstructured adaptive grids and employs multiple-scale finite differences to discretize the elliptic governing PDEs. These strongly coupled, highly nonlinear discretized equations are then solved simultaneously via Newton's method with a nested Bi-CGSTAB solver. The unstructured nature of the grids has necessitated substantial modification of the standard solver. Comparisons with computed results on equivalently refined conventional grids indicate that the LRR method provides substantial savings in execution time and computer memory requirements, with little or no loss in solution accuracy.

Three types of steady-state axisymmetric laminar methane/air flames are examined: a diffusion flame; a set of six partially premixed flames of various primary equivalence ratios; and a pair of lean and rich Bunsen flames. Each flame model employs a vorticity-velocity formulation of the governing equations and includes GRI-Mech chemistry, multicomponent transport, and optically thin radiation modelling. Agreement with available experimental measurements of temperature and species concentrations ranges from very good to excellent.

Speaker's web page: http://www.eng.yale.edu/faculty/vita/bennettbeth.html

Institution web page: http://www.yale.edu/

July 24, 2000

First-Order System LL* (FOSLL*): Scalar Elliptic Partial Differential Equations

Zhiqiang Cai
Purdue University

Email: zcai@math.purdue.edu

Abstract:

The *L^2*-norm version of first-order system least squares (FOSLS) attempts to reformulate a given system of partial differential equations so that applying a least-squares principle yields a functional whose bilinear part is *H^1*-elliptic. This means that the minimization process amounts to solving a loosely coupled system of elliptic scalar equations. An unfortunate limitation of the *L^2*-norm FOSLS approach is that this product *H^1* equivalence generally requires sufficient smoothness of the original problem. Inverse-norm FOSLS overcomes this limitation, but at a substantial loss of real efficiency. The FOSLL* approach introduced here is a promising alternative that is based on recasting the original problem as a minimization principle involving the adjoint equations. This paper provides a theoretical foundation for the FOSLL* methodology and illustrates its performance by applying it numerically to several different examples. Results for the so-called two-stage approach applied to discontinuous coefficient problems show efficiency that is much better than inverse-norm least-squares methods. FOSLL* appears to exhibit the generality of the inverse-norm FOSLS approach while retaining the full efficiency of the *L^2*-norm approach. This is a joint work with Manteuffel, McCormick, and Ruge at University of Colorado at Boulder.

Speaker's web page: http://www.math.purdue.edu/~zcai/

Institution web page: http://www.purdue.edu/

July 19, 2000

Multigrid for 2-D High Reynolds Number Turbulent Transonic Flows

Steven Allmaras
Boeing

Email: steven.r.allmaras@boeing.com

Abstract:

We present multigrid methods for the solution of the steady-state 2-D compressible Navier–Stokes equations discretized by an upwind finite-volume scheme on structured grids. A one-equation turbulence model is included in the set of equations. We describe the appropriate multigrid ingredients necessary to achieve effective convergence rates for high Reynolds number turbulent transonic flows. We describe trades for the formulation of the equations to be relaxed, the multigrid coarsening strategy, the choice of relaxation on each grid, and the degree of coupling between the conservation equations and the turbulence model equation. We also propose modifications for the one-equation turbulence model to alter its non-physical transient behavior as an aid to achieve steady-state solutions; these modifications, however, do not change the steady-state formulation of the model.

We present results for two cases: an inviscid transonic channel flow and a turbulent transonic airfoil. The achieved convergence rates are substantial improvements over those obtained by state-of-the-art methods.

We finish with a discussion of the extension of the method to 3-D structured grids and possibilities for extension to unstructured grids. Issues related to parallel implementation are also discussed.

Institution web page: http://www. boeing.com/

June 27, 2000

Liquid Crystalline Polymers and Other Macromolecular Fluids

Gregory Forest
University of North Carolina

Email: forest@amath.unc.edu

Abstract:

Many high-performance materials are made from macromolecular fluids, ranging from spider silk to synthetic liquid crystalline polymers like Kevlar and Vectra. Flows of these materials couple the dynamics of anisotropic macromolecules with hydrodynamics. There are many length-scales, and approaches, in the modeling of these systems. We have focused on a mesoscale description following Doi & Edwards, in terms of an orientation tensor for the second moment of the molecules with respect to a probability distribution function. This object relates directly to measurements of quantities such as birefringence or light scattering intensities. We will describe various phenomena we have studied and advertise some interesting challenges. Topics include a revisitation of the classical isotropic-to-nematic phase transition, what happens to this transition in the presence of imposed flows, and methods to construct heterogeneous patterns routinely observed in experiments. Collaborators include Qi Wang, IUPUI, and Hong Zhou, UC-Santa Cruz.

Speaker's web page: http://www.amath.unc.edu/Faculty/forest/

Research web page: http://www.amath.unc.edu/

Institution web page: http://www.unc.edu/

June 14, 2000

On the Analysis and Construction of Absorbing Layers in CEM and CFD

Saul Abarbanel
Tel Aviv University

Email: saula@math.lbl.gov

Abstract:

This talk will review some recent results about the theory and application of PML (Perfectly Matched Layers):

1. It will be shown that that set of split PDE's of the Berenger PML method is not strongly well-posed. Under certain perturbations its solutions may be inappropriate.

2. A recently introduced system of PDE's (Zhao & Cangellaris, 1996; Ziolkowski, 1997; Petropoulus et al., 1998), based on physical considerations, and which describes the behavior of EM waves in artificial absorbing layers, is analyzed. Analytic solutions are found for the cases of semi-infinite and finite depth layers, both for primitive and characteristic boundary conditions.

A different set of equations, based on mathematical considerations, and which offers some advantages, is constructed and analyzed for the same geometries and boundary conditions.

Both sets of PDE's are strongly well posed.

3. A set of strongly well-posed PML equations is developed for describing the absorption of acoustic and vorticity waves in two-dimensional convective aeroacoustics under the assuption of a constant mean flow. A central piece in the derivation is the development of a coordinate transformation that conserves the dispersion relation from the static acoutic case. Numerical solutions illustrate the efficacy of this PML scheme even when applied to the fully nonlinear Euler equations.

This work was done in collaboration with David Gottlieb and Jan Hesthaven of Brown University.

Institution web page: http://www.tau.ac.il/

June 12, 2000

Using a Multilayer Ocean Model on Clusters versus a Traditional Supercomputer

Craig Douglas
University of Kentucky

Email: douglas@ccs.uky.edu

Abstract:

We simulate oceanic overflow problems. The long-term goal of this research is to understand both the local dynamics of downslope flows in the ocean and their role in the Earth's global thermohaline circulation. The modeling of these flows and their climatic impact is complicated by the inherent range of spatial scales involved, which extend from the global scale [O(10,000) km] down to the local scale of the overflows themselves [O(1) km], and by the intrinsic three dimensionality of the overflow dynamics.

The Spectral Element Ocean Model (SEOM) offers an elegant solution to these difficulties. It features advanced algorithms, based on *h-p* type finite element methods, allowing accurate representation of complex coastline and oceanic bathymetry, variable lateral resolution, and high-order solution of the three dimensional oceanic equations of motion.

SEOM's geometrical flexibility permits highly inhomogeneous horizontal grids. An added advantage of the technique is its scalability. Most of the computations are carried out at the element level; only interface information needs to be exchanged between elements. The dual characteristic of dense and structured local computations, and sparse and unstructured communication, enhances the locality of the computations, and makes SEOM ideally suited for parallel computers.

In this talk we demonstrate what types of cluster machine characteristics are suitable for solving our problems, how much they cost at what point in time, and how they compare to a traditional RISC based supercomputer solution. Our comparisons will note time, money, and aggravation factors.

Speaker's web page: http://www.ccs.uky.edu/~douglas/

Research web page: http://www.mgnet.org/

Institution web page: http://www.uky.edu/

June 12, 2000

Random Testing with "Fuzz": A Decade of Random Testing

Barton Miller
University of Wisconsin

Email: bart@cs.wisc.edu

Abstract:

In 1990, we published the results of a study of the reliability of standard UNIX application/utility programs. This study showed that by using simple (almost simplistic) random testing techniques, we could crash or hang 25%–33% of the these utility programs. In 1995, we repeated and significantly extended this study using the same basic techniques: subjecting programs to random input streams. This study also included X-Window applications and servers. A distressingly large number of UNIX applications still crashed with our tests. The X-window applications were at least as unreliable as command-line applications. The commercial versions of UNIX fared slightly better than in 1990, but the biggest surprise was that Linux and GNU applications were significantly more reliable than the commercial versions.

This year (2000), we took another stab at random testing, this time testing applications running on Windows/NT . Given valid random mouse and keyboard input streams, we could crash or hang 45% of these applications.

In this talk, I will discuss our simple testing procedure as it applies to UNIX and Windows. I will then present the UNIX and NT test results. These results include, in many cases, identification of the bugs and the coding practices that caused the bugs. In several cases, these bugs introduce issues relating to system security. The talk will conclude with some philosophical musings on the current state of commercial software.

Papers on the three studies (1990, 1995, and 2000) can be found at:
ftp://grilled.cs.wisc.edu/technical_papers/fuzz.ps
ftp://grilled.cs.wisc.edu/technical_papers/fuzz-revisited.ps
ftp://grilled.cs.wisc.edu/technical_papers/fuzz-nt.ps

PDF versions are also available (change the .ps to .pdf). The tools used for the testing can be found at:
ftp://grilled.cs.wisc.edu/fuzz/

Institution web page: http://www.wisc.edu/

June 8, 2000

Sputnik: A Programming Model for Automated Decomposition on Heterogeneous

Sean Peisert
San Diego Supercomputer Center

Email: peisert@sdsc.edu

Abstract:

Clusters of multiprocessor nodes are becoming common in scientific computing. As a result of the expandability of clusters, faster nodes are frequently added and older nodes are gradually removed, making the cluster heterogeneous. As heterogeneity increases, traditional methods for programming clusters of multiprocessors become less optimal, because they do not account for the fact that a cluster will only run as fast as the slowest node. Sputnik is a programming methodology and software library that addresses the problem of heterogeneity on a dedicated cluster of multiprocessors.

Sputnik uses a two-stage process for running applications on a cluster of multiprocessors. The first stage assesses the relative performance of each node by running the program individually on each node, determining from the run times both the performance and application-specific optimization. Using the timings obtained from stage one, the second stage partitions the dataset non-uniformly, according to the relative speed of each node. All future runs of the program use the optimal partitionings and number of threads per node.

Sputnik is implemented on top of the KeLP infrastructure to handle irregular decomposition and data motion. It enables code to be written for a heterogeneous cluster as if the cluster is homogeneous. Sputnik can run scientific applications on a heterogeneous cluster faster, with improved utilization, than a nearly identical program written in KeLP alone. Experimental results from a pair of SGI Origin2000's indicate that Sputnik can improve run-time of an iterative solver for Poisson's equation by 35%.

Speaker's web page: http://www.sdsc.edu/~peisert/

Institution web page: http://www.sdsc.edu/

June 5, 2000

Development and Application of a Parallel, Solution-Adaptive Code for Space Physics

Kenneth Powell
University of Michigan

Abstract:

Over the past several years, a multidisciplinary group in space physics, CFD and computer science has worked to develop a solution-adaptive, parallel code for space physics. The uses dynamic adaptation to resolve the multiple scales inherent in this class of problems, and has achieved high parallel efficiency on a number of machines. On a 1490-processor Cray T3E, speeds above 300 GigaFLOPs have been sustained, even while using the dynamic adaptation abilities of the code. The development of the current code, some applications of the current code, and plans for future work will be discussed.

Speaker's web page: http://www-personal.engin.umich.edu/~powell/

Research web page: http://www.engin.umich.edu/dept/aero/people/faculty/powell.html

Institution web page: http://www.umich.edu/

May 31, 2000

Large Multiscale Tensor-Product Finite Element Approximation Spaces

Stephan Knapek
University of Bonn

Email: knapek1@llnl.gov

Abstract:

Despite the rapid development of computer technology, many problems in scientific computation remain intractable to a direct numerical simulation due to storage requirements and numerical complexity. Especially problems in higher dimension remain well out of the range of todays largest computers even after the use of parallelization. This intractability of higher-dimensional problems is often refered to as the "Curse of dimension."

The situation looks different when we impose specific regularity assumptions on our problem. Then, Finite Element subspaces with relatively large dimensions that contribute only "little" to the error reduction can be identified and omitted from the approximation space. This idea has appeared under various names (boolean blending schemes, hyperbolic cross points, sparse grids, ...) in approximation and interpolation theory and attracted much attention in recent time.

In this talk we generalize the concept of hyperbolic cross points and discuss applications to operator equations such as integral and partial differential equations. We construct a scale of nested finite element spaces based on multiscale tensor-product bases, and we show under which circumstances these approximation spaces break the curse of dimension.

Institution web page: http://www.uni-bonn.de/unibonn_en.html

May 25, 2000

MPI Software Technology's Path Forward Program for Terascale MPI Deployment on ASCI Architectures

Anthony Skjellum
MPI Software Technologies Inc.

Email: tony@mpi-softtech.com

Abstract:

In this presentation, we discuss the technology underpinnings and plans for achieving a new-generation of ASCI-relevant MPI implementations, based on the Software Pathforward Program. MPI Software Technology, Inc. (MSTI) is providing improved MPI-1 and MPI-2 capabilities for tri-lab ASCI systems, based on its existing and developing commercial technology for high performance messaging middleware. Specific technological achievements, including thread safety, high performance derived data types, and overlapping of communication and computation, are mentioned. This talk also provides background on the company, its middleware products/technologies, and work in clusters and embedded MPI deployments. Strategies for moving MPI/Pro—the commercial-grade MPI offered by MSTI—to ASCI platforms is included.

At the conclusion, input from users about deficiencies in existing ASCI and cluster MPI settings is invited, as is input about MPI-2 perceptions/interest/priorities moving forward.

Speaker's web page: http://www.mpi-softtech.com/company/president.html

Research web page: http://www.mpi-softtech.com/

May 18, 2000

High Resolution Central Schemes

Eitan Tadmor
University of California, Los Angeles

Email: tadmor@math.ucla.edu

Abstract:

We discuss recent developments of high-resolution schemes for the approximate solution of hyperbolic conservation laws, Hamilton–Jacobi (HJ) equations and related nonlinear problems.

We focus on non-oscillatory central schemes as prototype for Godunov-type projection methods. A variety of numerical experiments demonstrate that the proposed central schemes offer simple, robust, Riemann-solver-free "black box" solvers, while at the same time they retain the high-resolution content of the more expensive upwind schemes.

Among several issues to be discussed will be the following highlighted features:

Scalar equations. Variation and entropy stability estimates as well as multidimensional L1-bounds are presented. New convergence results based on both global and pointwise error estimates follow.

Systems of equations. Extension to systems is carried out by componentwise application of the scalar framework. It is in this context that the central schemes offer a remarkable advantage over the corresponding upwind framework.

Multidimensional problems. Since we bypass the need for (approximate) Riemann solvers, multidimensional problems are solved without dimensional splitting. In fact, the proposed class of central schemes is utilized for a variety of nonlinear transport equations, and in this context we demonstrate the construction and implementation of central schemes for HJ and incompressible Euler equations.

We describe recent developments of new high-resolution central schemes for general convection-diffusion equations.

We overview recent applications to various models, including incompressible flows, MHD equations, simulations of semi-conductors models, and more.

Speaker's web page: http://www.math.ucla.edu/~tadmor/

Research web page: http://www.math.ucla.edu/faculty/tadmor.html

Institution web page: http://www.ucla.edu/

May 18, 2000

UCLA's new NSF Institute for Pure and Applied Mathematics: Program and Collaborations with LLNL

Eitan Tadmor
University of California, Los Angeles

Email: tadmor@math.ucla.edu

Abstract:

UCLA will host one of NSF's three mathematics institutes as of September 2000. The first-year program will consist of units on Functional Genomics, Financial Mathematics, Oscillatory Integrals and Dispersive Equations, and Geometrically Based Motions. An overview of the institute and of opportunities for LLNL-IPAM collaborations will be presented.

Speaker's web page: http://www.math.ucla.edu/~tadmor/

Research web page: http://www.math.ucla.edu/faculty/tadmor.html

Institution web page: http://www.ucla.edu/

May 9, 2000

Sparse Pattern Selection Strategies for Robust Frobenius Norm Minimization

Iain Duff
CCLRC Rutherford Appleton Laboratory

Email: I.S.Duff@rl.ac.uk

Abstract:

We consider preconditioning strategies for the iterative solution of dense complex symmetric non-Hermitian systems arising in computational electromagnetics. We consider in particular sparse approximate inverse preconditioners that use a static nonzero pattern selection. The novelty of our approach comes from using a different nonzero pattern selection for the original matrix than for the preconditioner and from exploiting geometrical information from the underlying meshes in addition to using methods based on the magnitude of the entries.

The numerical and computational efficiency of the proposed preconditioners are illustrated on a set of model problems arising both from academic and from industrial applications.

The results of our numerical experiments suggest that the new strategies are viable approaches for the solution of large-scale electromagnetic problems using preconditioned Krylov methods. In particular, our strategies are applicable when fast multipole techniques are used for the matrix-vector product on parallel distributed memory computers.

Speaker's web page: http://www.cse.clrc.ac.uk/Person/I.S.Duff/

Institution web page: http://www.rl.ac.uk/

April 28, 2000

Numerical Simulation Methods for Electromagnetic Scattering Problems

Karl Warnick

University of Illinois, Urbana-Champaign

Email:
warnick@sunchew.ece.uiuc.edu

Abstract:

Electromagnetic analysis of radar targets, antennas, microwave circuits, digital systems, communications links, and optical devices is a significant challenge in the field of scientific computation. Due to the indefinite, wavelike character of these problems, efficient methods such as multigrid, which are highly succesful for static and electrically small problems, break down as electrical size increases. For surface integral equation or boundary element solvers, the Fast Multipole Method (FMM) has brought dramatic progress towards attaining O(1) steps per degree of freedom. Three outstanding problems remain: grid robustness, accuracy control, and increasing computational efficiency in the high-frequency asymptotic limit. This presentation will explore these open questions and will outline current research directions, including methods for integrating multigrid with FMM in the high frequency regime.

Research web page: http://www.ccem.uiuc.edu:80/reschew.html

Institution web page: http://www.uiuc.edu/

April 25, 2000

Scalable Adaptive Algorithms, Scalable Platforms

Gerhard Zumbusch
University of Bonn

Email: zumbusch@iam.uni-bonn.de

Abstract:

We present some highlights of the parallel computing activities of our research group. These include the construction of a parallel computing platform, namely our Parnass2 PC cluster as Europe's fastest cluster so far, and the development of several parallel multilevel and multigrid algorithms, especially in the presence of adaptive grid refinement, and the related parallel load-balancing issues.

The cluster Parnass2 at the Iinstitute for Applied Mathematics at the University of Bonn was designed and built by the group of researchers. The 144-processor cluster is listed on both the current and the previous TOP500 list of the world's supercomputers. It is listed there among only two further clusters, namely the larger CPlant at Sandia (Albuquerque) and the Avalon cluster at LANL, both in the United States. Hence, the Parnass2 cluster is probably the cheapest computer on the lists. Together with the superior parallel performance of the Linpack benchmark (even among commercial clusters) it shows best price-performance ratio. We will discuss and compare several issues that lead to the current design, including network performance and software and SMP computing nodes. The comparison both to the Cray T3E and to IBM SP series (and to ASCI blue pacific) are especially of interest.

The second part of the talk is devoted to algorithmic developments. Here we concentrate on parallel implementations of adaptive grid refinement and the related load balancing issue. We propose space-filling curve partition methods for locality both on the message-passing level and within the memory of a single node. We discuss iterative multilevel and multigrid implemenations on top of such grids with grid-independent convergence rates and we highlight how the partion method affects the complexity of the implementation and the performance of such a PDE solver. Furthermore, alternative discretization schemes based on adaptive wavlet methods and their impact on parallelization are discussed.

Speaker's web page: http://wissrech.iam.uni-bonn.de/people/zumbusch.html

Research web page: http://wissrech.iam.uni-bonn.de/

Institution web page: http://www.uni-bonn.de/unibonn_en.htm

April 13, 2000

Techniques Leading to Scalable Algorithms for Sparse Linear Systems

X. Sherry Li
Lawrence Berkeley Laboratory

Email: xiaoye@nersc.gov

Abstract:

We present some useful techniques for sparse direct and iterative solvers on large-scale parallel machines. Many classical algorithms are re-designed and new algorithms developed to enhance scalability. One example is sparse LU factorization with partial pivoting. It is hard to implement efficiently on distributed memory machines, because of its somewhat unpredictable way of generating fine-grained work and intermediate results at runtime. We have investigated the possibility of replacing (dynamic) partial pivoting by some other techniques to control the element growth during the elimination. These include statically pre-pivoting large elements to the diagonal and dynamically modifying the small diagonal pivot. These numerical techniques enable us to achieve a more scalable MPI implementation of the LU factorization.

Another example is ordering of the equations and variables of a sparse linear system. For a direct solver, the goal of ordering is usually to minimize the fills in the factored matrices. But the best fill-reducing ordering may not give the best parallel runtime. For a CG-like solver, the ordering has great impact on the efficiency of the matrix-vector product kernel. We will illustrate these by comparing several popular ordering strategies.

Speaker's web page: http://www.nersc.gov/~xiaoye/

Research web page: http://www.nersc.gov/

Institution web page: http://www.lbl.gov/

April 6, 2000

Probabilistic Clustering of Sequences, Trajectories, and other Non-Vector Data

Padhraic Smyth
University of California, Irvine

Email: smyth@ics.uci.edu

Abstract:

Clustering has a long history in exploratory data analysis. Searching for homogenous data clusters is a fundamental component of any endeavour to discover structure from data. This talk will begin with a brief overview of a general model-based probabilistic clustering framework. This approach, relatively well known in applied statistics, focuses on the use of mixture models as an underlying generative model for observed data. It also provides a relatively objective and sound methodology for determining the model which is closest to truth.

A particularly useful aspect of the probabilistic approach is the ability to generalize from clustering in vector-spaces to clustering sequences, trajectories, and other "dynamic" data that we commonly observe from individuals and/or systems. The main part of the talk will introduce a new and general framework for mixture-model clustering of such data. The probabilistic approach solves in an elegant and coherent manner the dual difficulties of (a) how to define distance metrics between observations (e.g., sequences) of different sizes, and (b) how to "weight" different individuals for whom we have different amounts of data. Illustrative applications include unsupervised learning of gestures from video data, clustering of individuals based on Web browsing behavior, modeling of gene expression data, modeling of cyclone trajectories, clustering locusts based on observed motor behavior, and clustering of medical patients based on histograms of red blood cells. A subset of these applications will be discussed, time permitting. The talk will conclude with a brief discussion of some apparently useful connections between mixture modeling in this context and Bayesian hierarchical models.

Speaker's web page: http://www.ics.uci.edu/~smyth/

Institution web page: http://www.uci.edu/

March 29, 2000

Large-Scale PDE-Constrained Optimization: Parallel Algorithms and Applications to Optimal Design, Optimal Control, and Parameter Estimation Problems

Omar Ghattas
Carnegie Mellon University

Email: oghattas@cs.cmu.edu

Abstract:

Very large scale PDE-constrained optimization problems arise naturally in many areas of engineering and science, often taking the form of optimal design, optimal control, or parameter estimation problems. The common denominator is a nonlinear optimization problem that is constrained by the PDEs that govern behavior of the physical system. Thus, solving the PDEs is just a subproblem associated with optimization, which can be orders of magnitude more challenging computationally. Despite its importance, little attention has been devoted to the design of parallel algorithms for PDE-constrained optimization. This is expected: it makes little sense to address the "inverse" problem until the "forward" problem is well understood. However, recent advances in parallel PDE solvers and the arrival of the teraflop computing era motivate the pursuit of very large-scale simulation-based optimization. Despite these advances, many PDE-constrained problems remain intractable with current optimization technology. To render them tractable, we need to develop parallel optimization algorithms that exploit the PDE nature of the constraints, scale to the millions of constraints and variables that arise upon discretization, and capitalize on emerging highly parallel supercomputers.

I will give an overview of the TAOS (Terascale Algorithms for Optimization of Simulations) Project at CMU, whose goals are to develop the enabling parallel numerical algorithms for large-scale simulation-based optimization, and to apply them to driving optimal control, optimal design, and parameter estimation problems in engineering and science. The talk will be illustrated with motivating applications in each of the three classes: optimal design of artificial heart devices, optimal boundary control of viscous flows, and inverse earthquake ground motion modeling. I will describe a family of parallel Newton–Krylov methods for solution of the optimality system of PDE-constrained optimization problems. These are essentially full-space Newton-Sequential Quadratic Programming (SQP) methods with preconditioning by reduced-space limited memory quasi-Newton SQP and approximate forward Jacobians. This combines the rapid convergence of Newton methods with the low per-iteration cost of approximate methods. I will present studies of parallel efficiency and scalability of a PETSc-based implementation for the problem of optimal control of a viscous incompressible fluid by suction/injection of fluid on its boundary. Numerical experiments for problems of size up to 1.5 million state variables and 50,000 control variables on up to 256 processors of a Cray T3E-900 yield encouraging results. I will also discuss a space-and-time-efficient adjoint algorithm for time dependent PDE-constrained optimization, with applications to inverse wave propagation.

The TAOS Project is jointly directed with Larry Biegler. The research on Newton–Krylov optimization methods and optimal flow control is joint work with Ph.D. student George Biros; inverse earthquake modeling with Ph.D. students Volkan Akcelik and Yiannis Epanomeritakis and collaborator Jacobo Bielak; and artificial heart design with Ph.D. student Ivan Malcevic and collaborators Jim Antaki and Greg Burgreen (University of Pittsburgh Medical Center).

Speaker's web page: http://www.cs.cmu.edu/~oghattas/omar.html

Institution web page: http://www.cmu.edu/

March 28, 2000

Stable Embedded Finite-Difference Time-Domain Methods For Maxwell Equation

Jan Hesthaven
Brown University

Email: Jan_Hesthaven@brown.edu

Abstract:

The last decade has seen a renewed interest in the development of flexible, simple, and accurate methods for solving Maxwell's equations. These efforts have partly been guided by technological developments requiring the ability to model broadband signals and their interaction with very complex configurations of materials. To enable such flexibility much effort has been centered around the development of time-domain methods in which Maxwell's equations are solved as a hyperbolic system.

The most popular time-domain method currently used in electromagnetic computations is based on central second-order finite differences in space and a Leapfrog scheme in time implemented on Cartesian grids, staggered in time and space, and the computational problem is simply embedded into this grid by assigning different material properties to the grid points. The advantages of this approach, known as the Yee scheme in computational electromagnetics, are its simplicity and robustness which make it the favorite tool for practitioners. However, this approach suffers from a number of significant problems such as a need to employ staircase approximations to curved boundaries and interfaces and an inability to enforce the correct electromagnetic boundary conditions across material interfaces.

In this presentation, we discuss a stable second-order Cartesian grid finite difference scheme for the solution of Maxwell's equations. The scheme employs a staggered grid in space while several choices are possible for the temporal integration. However, contrary to the Yee scheme, this new scheme represents the physical location of the material and metallic boundaries correctly to the order of the scheme, hence eliminating the problems caused by staircasing, and enforces the correct jump-conditions on the individual field components across material interfaces.

Accuracy and strict stability are discussed for the one-dimensional scheme and issues related to two- and three-dimensional problems are briefly mentioned. The analysis exposes that the effects of staircasing as well as a lack of properly enforced jump-conditions on the field components have a significant impact on the global accuracy. Indeed, the analysis reveals that for cases where a field component is discontinuous along a gridline, as happens for general two- and three-dimensional curvilinear material interfaces, the Yee scheme may exhibit local divergence and loss of global convergence.

This new formulation, the computational cost of which is primarily incurred during a preprocessing stage, also lends itself to a very efficient parallel implementation by allowing for the definition of 'ghost'-interfaces to minimize the communication costs.

To validate the analysis several one- and two-dimensional test cases are presented, showing an improvement of typically 1-2 orders of accuracy at little or no additional computational cost over the Yee scheme, which in general remains first order accurate at best.

The work discussed is done in collaboration with Dr. Adi Ditkowski (Brown University), Dr. Kim Dridi (Ris National Laboratory, Denmark), and Mr. Chun-Hau Teng (Brown University).

Speaker's web page: http://www.cfm.brown.edu/people/jansh/

Research web page: http://www.cfm.brown.edu/home.html

Institution web page: http://www.brown.edu/

March 27, 2000

High-Order Unstructured Grid Methods for Conservation Laws

Jan Hesthaven
Brown University

Email: Jan_Hesthaven@brown.edu

Abstract:

Increasing interest in the modeling of unsteady problems over very long times and in computationally very large domains requires that accurate and efficient high-order methods be developed for such problems. The geometric complexity, however, of problems of scientific and industrial interest makes such developments a very significant challenge.

The application of high-order methods, and in particular spectral methods, is traditionally restricted to domains that can be smoothly mapped to a unit cube, hence allowing for the construction of a well-behaved multi-dimensional approximation using tensorproducts. Unfortunately, as is well known, automated grid generation and adaptive meshing is greatly complicated by using only hexahedrals as the fundamental element.

Guided by these observations we shall discuss how to develop accurate high-order/spectral methods on triangles and tetrahedra in which the approximation is based upon truly multi-dimensional Lagrange interpolation polynomials much in the spirit of classical spectral collocation methods. To enable such a construction we first discuss the construction of families of nodal sets well suited for high-order interpolation on the *n*-simplex and show how to address this issue through the solution of problems in generalized electrostatics. Among other things we emphasize how one can exploit the symmetries in the nodal sets to factorize the discrete operators into a sparse format, hence allowing for a the fast evaluation of matrix-vector products to approximate spatial derivatives.

For the approximation of initial boundary value problems we shall discuss how to construct stable methods on almost general nodal distributions. The central problem is how to impose the boundary conditions in a stable way - an issue we resolve by showing how to imposing the boundary conditions only weakly through a penalty term. As we shall show, the resulting schemes contain, as a special case, the discontinuous Galerkin method.

To illustrate the efficacy of the proposed approach for the solution of conservation laws, we present examples of solving the three-dimensional Maxwell's equations for time-domain scattering from very complex bodies and the two-dimensional compressible Euler equations for two-fluid problems.

While the use of *n*-simplices ensures significant geometric flexibility, the nodal based approximation has a further advantage in that it enables the construction of very efficient preconditioners by allowing for the use of finite element-based preconditioning. We shall illustrate this latter point by examples of solving the two-dimensional incompressible Navier-Stokes equations emphasizing the need to consider conditioning techniques to remove the *h* as well as the *p* dependency. We shall offer some suggestions on how to accomplish this within the proposed framework.

The work discussed is done in collaboration with Dr. Tim Warburton (Brown University) and Dr. Luca Pavarino (University of Milan, Italy).

Speaker's web page: http://www.cfm.brown.edu/people/jansh/

Research web page: http://www.cfm.brown.edu/home.html

Institution web page: http://www.brown.edu/

March 24, 2000

Charon Parallelization Toolkit

Rob Van Der Wijngaart

MRJ Technology Solutions, NASA Ames Research Center

Email: wijngaar@nas.nasa.gov

Abstract:

Message passing is among the most efficient methods for parallelization of scientific programs. Other strategies are available, but these often fail to deliver the flexibility, scalability, and high performance required for large-scale computations. This is especially true for applications developed for distributed-memory machines. Unfortunately, writing message-passing programs is a slow, cumbersome, and error-prone process.

Charon is an application-level library (C and Fortran) conceived to help ease the burden of the message-passing implementation of a class of 'difficult' parallel programs. This class is characterized by implicit numerical algorithms on structured grids that are not—or not easily—amenable to data-parallel solution techniques.

Charon supports piecemeal conversion of a serial design or legacy code to a fully parallel message-passing code by decoupling data distribution and parallel execution. This is achieved by tools that mimic serial program execution on distributed data. Through a sequence of tuning steps, again using tools provided by Charon, performance is gradually improved.

In this seminar the basics of the library will be discussed, and some examples of its use and performance will be presented.

Institution web page: http://www.arc.nasa.gov/

March 23, 2000

Programming Tools for Very Large Dataset Subsetting and Aggregation

Joel Saltz
University of Maryland

Email: saltz@cs.umd.edu

Abstract:

Increasingly powerful computers, clusters, and multiprocessor machines have enabled computational scientists and engineers to model physical phenomena in great detail. As a result, overwhelming amounts of data are being generated by scientific and engineering simulations. In addition, large amounts of data are being generated by sensors of various sorts, attached to devices such as satellites and microscopes. The exploration and analysis of the resulting large datasets plays an increasingly important part in many domains of scientific research. In this presentation we describe the design and development of software systems designed to address the need to subset, explore, analyze, process, and visualize very large datasets.

The first software system, the Active Data Repository (ADR), targets large disk-based datasets in processing environments with multiple processors and multiple disks. ADR is a C++ class library designed to support the development of servers that support multidimensional range queries with user defined aggregation and filtering functions.

The second, closely related, software system is a set of middleware infrastructure, called DataCutter, that supports the processing of scientific data collections stored in archival storage systems across a wide-area network. DataCutter also provides support for subsetting of datasets through multi-dimensional range queries along with support for invoking a sequence of user-defined filtering and aggregation functions. Processing, network, and data copying overheads are minimized by the ability to place filtering and aggregation functions on different platforms.

We will also discuss a compiler infrastructure that allow users to target ADR and DataCutter by specifying data distributions using XML and by writing user defined aggregation functions in a high-level language.

Institution web page: http://www.umd.edu/

March 16, 2000

Data Mining: An Overview

N. Radhakrishnan and Raju Namburu

U.S. Army Research Laboratory (ARL)

Abstract:

This talk will cover, in a light-hearted way, the emergence of data mining (a subset of Knowledge Discovery) in academia, industry, and the government. Questions like, what is data mining?, why is it needed?, how is it done?, etc., will be outlined. Several successful applications of data mining in industry and the government will be covered with special emphasis on scientific data mining. The talk will also highlight data mining advantages and pitfalls and the need for responsible management of data.

N. Radhakrishnan is the Director of the Corporate Information and Computing Directorate, and R. Namburu is the DOD thrust area lead for computational mechanics.

Research web page: http://www.arl.mil/cicd/

Speaker's web page: http://www.arl.mil/cicd/radha.html

Institution web page: http://www.arl.mil/

March 14, 2000

Programming Environments for Multi-Application Simulation

Eric DeSturler

University of Illinois, Urbana-Champaign

Email: sturler@uiuc.edu

Abstract:

The Center for Simulation of Advanced Rockets (CSAR) aims to develop simulation programs for the detailed, whole-system simulation of solid propellant rockets under normal and abnormal operating conditions. For the simulation of solid propellant rockets we need to solve highly complicated multi-physics problems, involving among others, computational fluid dynamics with particles and chemical reactions, turbulence, solid dynamics and combustion, and crack propagation.

Rather then incorporate the equations of the constituent problems in a single system and combine the technology of currently separate programs in a single program, we aim to continue to use the existing programs for each constituent part. We will solve each application separately. The mutual influence of connected structures and/or coupled physical and chemical processes is accounted for explicitly by coupling the simulations of all relevant structures and processes at their "boundaries." This preserves the efforts in existing programs and it allows the developer to make or keep the optimal choices (model, spatial discretization, time integration, and solution procedures) for each individual simulated process. The complication is, of course, how to couple the simulated processes in a physically and mathematically correct way and in a computationally efficient way, and how to run the various programs in concord. We aim to develop a programming environment that will support combining with minimal changes and in a short time frame multiple grid-based applications for complex multi-physics simulations, especially involving highly dynamic adaptive simulations. Such a programming environment will be useful for a wide range of applications besides rocket simulation.

I will give an overview of the work at the CSAR, our current efforts, and future plans for the programming environment.

Institution web page: http://www.uiuc.edu/

March 13, 2000

Source Code Engineering Using SNiFF+ for Large-Scale Software Development

Gil Chita
TakeFive Software

Email: gil@takefive.com

Abstract:

SNiFF+ is an interactive source code analysis tool. Gil Chita, a Field Application Engineer for Wind River Systems, Inc., will show how to use SNiFF+ to understand the workings of existing, unfamiliar code. The presentation will have three parts: a 10-minute slide show introduction and overview, a 40-minute demonstration of SNiFF+, and a 10-minute question and answer session. The presentation is targeted at engineers and scientists that need to write new software modules, integrate the modules into existing systems, and port them to different operating systems.

Institution web page: http://www.takefive.com/

March 7, 2000

Performance Instrumentation and Visualization with SvPablo

Luiz De Rose
IBM TJ Watson Research Center

Email: laderose@us.ibm.com

Abstract:

In this talk I will present the Source View Pablo, a language and architecture independent toolkit for performance instrumentation, visualization, and analysis of programs written in C, Fortran 77, Fortran 90, and HPF, on sequential and parallel systems. Besides capturing application data via software instrumentation, SvPablo exploits multi-platform hardware performance counters, via its interface with PAPI, in order to capture the interaction of software and hardware. During execution of the instrumented code, the SvPablo library captures data and computes performance metrics on the execution dynamics of each instrumented construct on each processor. Because only statistics, rather than detailed event traces, are maintained, SvPablo can handle measurements of programs that execute for hours or days on hundreds of processors. In addition, SvPablo can be easily extended to new contexts with minimal changes to the software infrastructure.

Research web page: www.research.ibm.com/actc/

March 3, 2000

Modeling the Performance of Sn Transport on Parallel Computers

Harvey Wasserman
Los Alamos National Laboratory

Email: hjw@lanl.gov

Abstract:

Particle transport via deterministic approaches such as the Discrete Ordinates Sn consumes enormous amounts of time during simulations of interest to ASCI. Moreover, the proportion of time spent in this application is expected to grow as ASCI begins to approach its canonical goal of order billion-cell problems. This talk will describe attempts to create a predictive capability for performance of a Cartesian-grid Sn transport application. The performance model attempts to parameterize characteristics of the machine (such as message latency and bandwidth and a characteristic floating-point computation rate) and of the application (such as number of "sweeps" and blocking factors for non-parallel computation). The talk will also describe attempts to generalize the model from a uniform network topology to lower-dimensional networks such as those used in clusters of SMPs.

Speaker's web page: http://www.c3.lanl.gov/~hjw/hjw.html

Research web page: http://www.c3.lanl.gov/cic19/teams/par_arch/par_arch.shtml

Institution web page: http://www.lanl.gov/

March 2, 2000

Hybrid MPI-OpenMP Implementation of an Unstructured Grid Agglomeration Multigrid Solver

Dimitri Mavriplis
ICASE NASA Langley Research Center

Email: dimitri@icase.edu

Abstract:

A two-level MPI-OpenMP parallel implementation of an existing unstructured multigrid Navier-Stokes solver is described. For both programming models, parallelism is achieved through domain partitioning applied to the various multigrid levels. In the MPI programming model, each process operates on a single partition and communication between the partitions is handled by explicit message-passing routines. In the OpenMP model, loops over the number of domains are parallelized using compiler directives, thus spawning one thread for each partition, and communication is handled by copying values from ghost vertices in individual partitions to their real images in adjacent partitions. The current implementation allows the solver to be executed either in a pure MPI mode, a pure OpenMP mode, or a two-level MPI-OpenMP mode. In the latter case, two hybrid communication strategies for combining MPI and OpenMP have been implemented. In the first approach, MPI messages are sent and received by individual OpenMP threads running in parallel. In the second approach, the thread-level messages are packed into larger messages which are then sent to remote MPI processes using the master thread, where they are subsequently unpacked.

Performance comparisons of MPI, OpenMP, and combined MPI/OpenMP using both communication strategies are given on the Cray-SV1, and the SGI Origin-2000. Benchmark results are demonstrated on all three current open ASCI machines using up to 1024 and 2048 processors.

The two-level parallel unstructured multigrid Navier–Stokes solver described herein will form the basis of a radiation diffusion solver, which is currently under development in two dimensions.

Institution web page: http://www.larc.nasa.gov/

February 28, 2000

SCM by Perforce

Ines Heinz
Lawrence Livermore National Laboratory

Email: heinz2@llnl.gov

Abstract:

Perforce is a software configuration management system that was designed to provide a great deal of control while presenting a minimal footprint. The architecture of Perforce will be presented along with a brief comparison to ClearCase and CVS. Some strengths and weaknesses will be outlined. Finally, there will be a brief discussion of features and personal experiences.

Institution web page: http://www.llnl.gov/

February 24, 2000

Fast Solvers in Scientific Computing

Gabriel Wittum
University of Heidelberg

Email: wittum@iwr.uni-heidelberg.de

Abstract:

Solvers play a crucial role in many simulations determining the complexity of the overall process. Thus, solving often limits the obtainable accuracy and fast solvers are the key to simulating new challenging problems.

In this lecture, numerical strategies for the simulation of large problems are discussed. The main directions are adaptivity, multigrid, and parallelism. The software system UG is presented which is based on these strategies. In several application cases the efficiency of the selected approach is shown.

Research web page: http://www.iwr.uni-heidelberg.de/~techsim/

Institution web page: http://www.uni-heidelberg.de/index_e.html

February 23, 2000

An Asymptotic Analysis of Multigroup Neutron Diffusion

Gerald Hedstrom
LLNL Retiree

Email: ghedstrom@home.com

Abstract:

It has been known for many years that for monoenergetic neutrons the diffusion equation is an asymptotic limit of the Boltzmann equation, and this information is used in the design of proconditioners for numerical methods. In this talk, I use the theory of positive operators to extend the results to the energy-dependent case.

February 11, 2000

Compiler Technology and Scientific Computing

David Padua
University of Illinois

Email: padua@uiuc.edu

Abstract:

Scientific computing has always been an important focus for compiler technology, to the point that the first commercial compiler was a Fortran compiler and many of the early optimization techniques were developed to improve the performance loops containing array accesses. Today, powerful compilers are an important part of any development environment for numerical codes. These compilers usually generate very good code that can take advantage of the capabilities of the most powerful processors.

Although much has been achieved since the first Fortran compiler, there are still many interesting and important issues in compiler technology that need attention. For example, techniques to analyze accesses to irregular data structures are not well understood although they are of great importance for conventional optimizations as well as detection of instruction level parallelism. Effective algorithms to handle explicitly parallel programs are also important because parallel programming is likely to become more frequent as SMPs grow in popularity. Today's techniques were developed for sequential programs and, therefore, do not always work correctly in a parallel context. For example, some compilers eliminate as dead code spin-lock loops, which thereby produces incorrect code and creates difficulties for the unsuspecting programmer.

Also, techniques to detect coarse grain parallelism are quite important and more needs to be done in this area, although the problem has proven much more difficult than originally expected. Other two issues are the systematization of the optimization process, which traditionally has been based on ad-hoc procedures that are not well understood, and the development of compiler algorithms that take advantage of information available in very high level programs to perform advanced optimizations.

In this talk, I will present a brief description of ongoing work at Illinois on these five topics.

Speaker's web page: http://www.cs.uiuc.edu/contacts/faculty/padua.html

Research web page: http://polaris.cs.uiuc.edu/polaris/polaris.html

Institution web page: http://www.u iuc.edu/

January 28, 2000

Incremental Nearest Neighbor Finding

Hanan Samet
University of Maryland, College Park

Email: hjs@umiacs.umd.edu

Abstract:

An algorithm for finding nearest neighbors incrementally is presented. It finds use in ranking spatial objects according to increasing distance from a query object in a graphical database. The algorithm makes use of a hierarchical spatial data structure. The algorithm is general in the sense that it is applicable to all hierarchical spatial data structures, all spatial data types, and to data of arbitrary dimensionality. The intended application area is a database environment, where the spatial data structure serves as an index. The algorithm is incremental in the sense that objects are reported one by one, so that a query processor can use the algorithm in a pipelined fashion for complex queries involving proximity. It is well suited for *k* nearest neighbor queries, and has the property that *k* needs not be fixed in advance. The algorithm has been used as the basis of a browser for graphical objects in a relational database. It also serves as the heart of the VASCO spatial data structure applet which awill be presented in the talk (found at *http://www.cs.umd.edu/~hjs/quadtree/index.html*). Other applications of the algorithm include a distance semi-join which enables the computation of a discrete Voronoi diagram in an incremental manner. This is the result of joint work with Gisli R. Hjaltason.

Speaker's web page: http://www.cs.umd.edu/~hjs/

Institution web page: http://www.umcp.umd.edu/

January 27, 2000

Chemometric Analysis of ICP-AES Spectroscopic Data Collected on Acidic Aqueous Samples Mimicking Nist Plain Carbon Steels

John Quagliano
Los Alamos National Laboratory

Email: JOHNQ-MAN@prodigy.net

Abstract:

There is considerable interest in using Inductively Coupled Plasma Atomic Emission Spectroscopy to characterize the minor elemental components of stainless steel. Iron interferes spectrally with many of the emission lines of the other analytes. Tedious and laborious wet chemical separations (to remove Fe) were previously required to obtain quantitative determinations. The present work demonstrates the application of a multivariate statistical analysis for the determination of six transition metals in a high concentration aqueous iron matrix. Sample preparation involved only a simple nitric acid dissolution. A 60 sample calibration set with a solvent background subtraction yielded a Partial Least Squares (PLS) model with 9 principal components after a leave-one-out cross validation. After independent validation, all analytes but vanadium were predicted with errors of no worse than 11%. Based on this preliminary investigation, we have learned new implementation strategies that should enable us to improve the quality of our results.

Institution web page: http://www.lanl.gov/

January 27, 2000

Demonstration of Improved Analytical Chemical Operations Using the Analyst-Interactive Program FRAD (Fast Reduction of Analytical Data) Running Under Windows 95/98/NT an Application to ICP Mass Spectrometry

John Quagliano
Los Alamos National Laboratory

Email: JOHNQ-MAN@prodigy.net

Abstract:

Modern analytical instrumentation is capable of producing large data sets in a short amount of time due in part to multi-channel detection and computer automation. However, it is still the analyst's responsibility to interpret and organize the raw data into a useful and meaningful format and also to assure that numerous Quality Control (QC) criteria are satisfied. It has been our experience that instrument vendor and third party software programs are inadequate for these data reduction needs. Consequently, our chemists spend an inordinate amount of time performing spreadsheet manipulations to complete the data reduction. Therefore, there is much less time spent in the lab developing improved measurement techniques and the analytical services team produces a slower turnaround of results to customers.

To solve this problem, we have developed a stand-alone executable software program that formats raw data, performs QC checks, notifies the analyst of errors, prepares data for LIMS, and generates customer deliverable reports. The program FRAD is mouse/menu driven with the familiar Windows Graphical User Interfaces (GUIs) and is suitable for both technicians and scientists or engineers. As an example, this presentation will demonstrate how a 6-isotope ICP-MS raw data file that previously required 2 hours to process should now take less than half an hour to process without any loss of thoroughness or quality.

Our ultimate goal is to have this conveniently formatted set of data results serve as input for the commonly used and already available chemometric software packages. In this way, the combined FRAD/Chemometrics software can (a) make the basic chemometric tools more attractive to analytical service personnel and (b) assist R&D personnel in assessing the quality of their research data before treatment by advanced chemometric tools.

Institution web page: http://www.lanl.gov/

January 24, 2000

Parallel Program Analysis Framework for the DOE ACTS Toolkit

Allen Malony
University of Oregon

Email: malony@cs.uoregon.edu

Abstract:

The diversity of high-performance computing platforms, programming systems, and applications and libraries that make up the computing environments at the DOE laboratories requires a rich set of tools. It is difficult, however, to build tools that are both portable across the set hardware and software systems, and able to be integrated with other computing and program components. This talk describes our work in developing program analysis technology for static and dynamic analysis of parallel programs. The work is funded by the DOE 2000 program and targets the analysis requirements of the ACTS toolkit.

Three project activities are covered: TAU, PDT, and the TAU distributed monitor. The TAU (Tuning Analysis Utilities) activity is developing a robust performance analysis framework that has been demonstrated on all ASCI platform types and across all ACTS software layers. The PDT (Program Database Toolkit) activity is developing a code analysis system that can be used to build source-level tools. Finally, the TAU distributed monitoring system is enabling runtime access to TAU performance data.

This work will be described in detail and future activities will be discussed.

Research web page: http://www.cs.uoregon.edu/research/paracomp/proj/tau/

Institution web page: http://www.uoregon.edu/

January 19, 2000

Support for Dynamic Thread Management in Clustered Systems

Steve Chapin
Syracuse University

Email: chapin@ecs.syr.edu

Abstract:

The Heterogeneous Operating System Software (HOSS) project is building operating system software for high-performance, heterogeneous clusters. Although these clustered environments are capable of sustaining computation rates that rival or surpass conventional supercomputers, the performance delivered to scientific applications is only a fraction of this maximum. The intent of this research is to develop mechanisms for enabling high-level software such as applications, message-passing libraries, and runtime systems to better manage resources through the operating system kernel. The resulting operating system will provide management functions allowing middleware to predict the performance ramifications of policy choices, and to allow runtime systems to make quality-of-service guarantees to application programs.

Speaker's web page: http://www.cs.virginia.edu/~chapin/

Institution web page: http://www.syr.edu/

January 14, 2000

An Architecture for Management and High-Performance Access of Data

Alok Choudhary
Northwestern University

Email: choudhar@ece.nwu.edu

Abstract:

Large-scale scientific applications (e.g., ASCI applications) consume and produce tremendous amounts of data for analysis, visualization, and other purposes. It can be an overwhelming task just to manage this data for a large number of experiments. This problem is further exacerbated when performance is critical (in particular I/O), sharing data among multiple applications is required, and hierarchical and distributed storage systems are used. The number of system and performance-related parameters can be large, thereby making performance optimizations for I/O a daunting task where a user has to understand what optimizations are suitable under what conditions and how to achieve them.

In this talk we will present an architecture and initial implementation for a system for managing scientific data with a particular emphasis on "performance management" and "automatic performance optimizations" for data access. A DBMS is used to manage performance metadata which can store and use data access patterns and storage patterns, but data access is done using a runtime system outside the DBMS. In particular we will present (1) the architecture and design philosophy, (2) the runtime system and automatic access optimizations using the database model, (3) initial performance results, (4) initial results for HPSS access patterns, and (5) user interface for management using runtime system. We will also summarize other work as part of out project.

Speaker's web page: http://web.ece.nwu.edu/~choudhar/

Institution web page: http://www.nwu.edu/

January 13, 2000

The PARASOL Project: Overview and Domain Decomposition Solvers

Petter Bjorstad
University of Bergen, Norway

Email: petter@ii.uib.no

Abstract:

Part I: An overview of the PARASOL sparse linear solvers. The Parasol project, an R&D project supported by the European Commission from 1996 through the summer of 1999, will be described, with a special focus on the sparse, linear solver developments.

Part II: Experiences with Domain Decomposition Algorithms in the PARASOL project. This talk will describe the effort to create a framework for Domain Decomposition algorithms (or preconditioners) within the PARASOL project, a large R&D project on industrial strength, parallel, and sparse linear solvers.

The main emphasis will be on issues related to a robust implementation of the Neumann–Neumann algorithm for realistic problems resulting from finite element discretization. We will discuss the construction of appropriate coarse spaces and several issues related to hybrid sparse algorithms, that is, the combination of direct and iterative sparse methods.

Speaker's web page: http://www.ii.uib.no/~petter/

Research web page: http://www.parallab.uib.no/

January 11, 2000

Very High Resolution Simulation of Compressible Turbulence on the IBM-SP System

Art Mirin

Lawrence Livermore National Laboratory

Winner of 1999 Gordon Bell Award for Performance

A. Mirin, R. Cohen, B. Curtis, W. Dannevik, A. Dimits, M. Duchaineau, D. Eliason and D. Schikore, LLNL, S. Anderson, D. Porter and P. Woodward, University of Minnesota, L. Shieh and S. White, IBM

Abstract:

Understanding turbulence and mix in compressible flows is of fundamental importance to real-world applications such as chemical combustion and supernova evolution. The ability to run in three dimensions and at very high resolution is required for the simulation to accurately represent the interaction of the various length scales, and consequently, the reactivity of the intermixing species. Toward this end, we have carried out a very high resolution (over 8 billion zones) 3-D simulation of the Richtmyer-Meshkov instability and turbulent mixing on the IBM Sustained Stewardship TeraOp (SST) system, developed under the auspices of the Department of Energy (DOE) Accelerated Strategic Computing Initiative (ASCI) and located at Lawrence Livermore National Laboratory.

We have also undertaken an even higher resolution proof-of-principle calculation (over 24 billion zones) on 5832 processors of the IBM system, which executed for over an hour at a sustained rate of 1.05 Tflop/s, as well as a short calculation with a modified algorithm that achieved a sustained rate of 1.18 Tflop/s. The full production scientific simulation, using a further modified algorithm, ran for 27,000 timesteps in slightly over a week of wall time using 3840 processors of the IBM system, clocking a sustained throughput of roughly 0.6 teraflop per second (32-bit arithmetic). Nearly 300,000 graphics files comprising over three terabytes of data were produced and post-processed.

The capability of running in 3-D at high resolution enabled us to get a more accurate and detailed picture of the fluid-flow structure—in particular, to simulate the development of fine-scale structures from the interactions of long- and short-wavelength phenomena, to elucidate differences between two-dimensional and three-dimensional turbulence, to explore a conjecture regarding the transition from unstable flow to fully developed turbulence with increasing Reynolds number, and to ascertain convergence of the computed solution with respect to mesh resolution.

Speaker's web page: http://www.llnl.gov/CASC/people/mirin/

Project web page: http://www.llnl.gov/CASC/asciturb/

January 6, 2000

Parallel Implementation of Multigrid Methods on Unstructured Grids using Adaptive Finite Element

Linda Stals
Old Dominion University

Email: stals@cs.odu.edu

Abstract:

Solving systems of equations using the multigrid method on a parallel machine presents a challenge due to the need to handle both the intra-grid and inter-grid data dependencies. Our approach to tackling such complicated data dependencies is to implement a very flexible node-edge data structure using an object-oriented framework. This setup has allowed us to write the code in a manner which more closely mimics our conceptual view of the problem and thus we are able to refine the grids (in 2D and 3D), build the system of equations, balance the load, and solve the system of equations all in parallel.

During the presentation I plan to present examples from the solution of the radiation transport equations (ODU and ICASE). I shall also briefly mention some other applications, including the Plasma Ion Immersion Process (Universitat Augsburg) and flow in heterogeneous media (University of Bath).

Speaker's web page: http://www.icase.edu/~stals/

Research web page: http://www.icase.edu/~keyes/asci/

Institution web page: http://web.odu.edu/

December 14, 1999

Scheduling and Processor Allocation for Matrix Chain Products on Parallel Systems

Jong Kim

Pohang University of Science and Technology

Email: jkim@postech.ac.kr

Abstract:

Matrix computations are fundamentals of scientific computing and commonly arising computationally intensive parts of many numerical and non-numerical applications. In this talk, I present the computation of matrix chain products on parallel systems.

The problem we are considering is finding an optimal product schedule for evaluating a chain of matrix products on a parallel computer (the matrix chain scheduling problem, MCSP). Since the approach of parallelizing each matrix product after finding an optimal product sequence for single processor systems does not always guarantee the minimum evaluation time on parallel systems, we introduce a new processor scheduling and processor allocation algorithms for the MCSP which reduces the evaluation time of a chain of matrix products on a parallel computer. We show that the proposed scheduling algorithm significantly decreases the time required to evaluate a chain of matrix products in parallel systems. The scheduling approach may also improve the performance of many other scientific computations.

Speaker's web page: http://www.postech.ac.kr/~jkim/

Institution web page: http://www.postech.ac.kr/e/

December 13, 1999

The Urgent Need for Configuration Management

Michael Donaldson
Merant

Email:
Michael.Donaldson@merant.com

Abstract:

Configuration management—the disciplined approach to managing the evolution of software development and maintenance practices and their products—has matured to the point where organizations now cannot thrive without configuration management (CM) capabilities. Industry has learned many lessons from its software crisis: namely, that CM is a mission-critical, foundational practice. Organizations have seen the value of a full-featured process-oriented CM solution. They include:

- Improved release cycle times
- Reduction in bugs
- Automatic quality control
- Minimal change complexit
- Variant/concurrent release management
- Visibility into all work status
- Better forecasting of release dates
- Increased preparedness for audits
- Repeatability of processes and product creation
- Improved responsiveness to customers
- Faster standards certifications

PVCS Dimensions is a full-featured, process-oriented CM solution that provides systemic quality automation by operationally supporting: version control, parallel and distributed development, build and release management, workspace and repository management, change and issue tracking, status reporting, and audit control.

This presentation discusses why CM is urgently needed today, the role CM plays in process improvement goals, and how PVCS Dimensions maximizes your investment in software development.

Institution web page: http://www.merant.com/

December 3, 1999

Smarter Memory Controllers: Improving Memory System Performance from the Bottom Up

Sally McKee
University of Utah

Email: sam@cs.utah.edu

Abstract:

Microprocessor speed is increasing much faster than memory system speed: both speed-growth curves are exponential, but they represent diverging exponentials. The traditional approach to attacking the memory system bottleneck has been to build deeper and more complex cache hierarchies. Although caching may work well for parts of programs that exhibit high locality, many important commercial and scientific workloads lack the locality of reference that makes caching effective. A 1996 study by Sites and Perl on a commercial database workload shows that memory bus and DRAM latencies cause an 8X slowdown from peak performance to actual performance. Another 1996 study by Burger, Goodman, and Kagi finds that when compared to an optimal cache the efficiency of current caching techniques is generally less than 20% and that cache sizes are up to 2000 times larger. The evidence is clear: no matter how hard we push it, traditional caching cannot bridge the growing processor-memory performance gap.

This talk presents research that attacks the memory problem at a different level—the memory controller. We describe the Stream Memory Controller (SMC) system built at the University of Virginia, and the Impulse Adaptable Memory System being built at the University of Utah. The SMC dynamically reorders accesses to stream elements to avoid bank conflicts and bus turnaround delays and to exploit locality of reference within the page buffers of the DRAMs. Impulse virtualizes unused physical addresses and dynamically remaps data to improve cache and bus utilization. For instance, Impulse can gather sparse data into dense cache lines that have high locality and leave a smaller cache footprint. Gathering strided data improves performance on regular applications, and gathering data through an indirection vector improves performance on irregular applications (e.g., Impulse improves the NAS conjugate gradient benchmark's performance by 67%). By prefetching data within the memory controller, Impulse uses bus bandwidth more efficiently, transferring only the needed data when it is requested.

Both these systems require that the compiler or application writer supply the access pattern information that the memory controller exploits. The SMC optimizes low-level memory performance for regular applications. In contrast, Impulse optimizes performance within the cache hierarchy and the memory back end for both regular and irregular computations.

Speaker's web page: http://www.cs.utah.edu/~sam/

Institution web page: http://www.utah.edu/

December 2, 1999

Scalable Network Attached Secure Storage Systems

Garth Gibson
Carnegie Mellon University

Email: garth@cs.cmu.edu

Abstract:

While it is possible to construct off-the-shelf, widely distributed, and massively parallel storage systems with inherent high bandwidth, achieving low-latency file access remains a significant challenge. We are designing, implementing, and evaluating scalable, distributed, and parallel storage architectures, interfaces, and protocols to reduce access latency comprehensively. Our goal is to define the evolutionary path and revolutionary changes that will enable commodity storage components to be the building blocks of high-bandwidth, low-latency, secure scalable storage systems.

Our current definition for a NASD device includes all storage systems that exhibit direct client-drive data transfer in a networked environment and asynchronous oversight by the high-level filesystem.

Institutional web page: http://www.cmu.edu/

November 23, 1999

Quality Issues in Image Compression: Perceptual Measures and Processing for Performance Improvement

V. Ralph Algazi
University of California, Davis

Email: algazi@ece.ucdavis.edu

Abstract:

Two broad issues that determine the quality of compressed images are the perceptual importance of details or noise in the original image, and the artifacts introduced by the compressor. In this presentation, an overview and discussion of subjective and objective quality measures in the compression of achromatic still images are presented. We first consider subjective quality scales and measures based on properties of human perception, and then present and discuss some objective measures and their correspondence to the subjective assessment of quality. An interesting alternative to the compression of the original image is to process the image prior to compression, so as to reduce unimportant details while improving the compressibility of the image data. Some techniques for improving the bit-rate versus distortion trade-off by image processing are presented, as well as the quality measures that allow to quantify the performance improvement possible.

Speaker's web page: http://info.cipic.ucdavis.edu/~algazi/index.html

Institution web page: http://www.ucdavis.edu/

November 22, 1999

Generic Software Components for Scientific Computing

Jeremy Siek
University of Notre Dame

Email: jsiek@lsc.nd.edu

Abstract:

Developing high-quality software is an exceedingly difficult task. Although software engineers continue to develop new methodologies that promise to enable the development of high-quality software, many of these methodologies do not even live up to their own promises, much less to the more stringent requirements of scientific applications. All is not lost, however. Careful study of various software development paradigms in the context of hiqh-quality scientific software clearly points to several new approaches that demonstrate significant promise both in theory and in practice. We validate our approach with the recently developed Matrix Template Library (MTL), a comprehensive collection of generic software components for high-performance scientific computing. Four aspects of the design and development of MTL will be presented. First, we discuss the generic programming methodology used to create the MTL framework. Second, we describe the framework of the abstract MTL interface. Third, we present the concrete implementation of MTL. Finally, we give performance results demonstrating that the generic components MTL are able to achieve performance equal to that of vendor-tuned (and automatically tuned) numerical libraries.

Speaker's web page: http://www.lsc.nd.edu/~jsiek/

Research web page: http://www.lsc.nd.edu/research/mtl/

Institution web page: http://www.nd.edu/

November 22, 1999

The Generic Graph Component Library

Andrew Lumsdaine
University of Notre Dame

Email: lums@lsc.nd.edu

Abstract:

We present the Generic Graph Component Library (GGCL), a generic programming framework for graph data structures and graph algorithms. Following the theme of the Standard Template Library (STL), the graph algorithms in GGCL do not depend on the particular data structures upon which they operate, so that a single algorithm can operate on arbitrary concrete representations of graphs. To attain this type of flexibility for graph data structures, which are more complicated than the containers in STL, we introduce several important concepts that form the generic interface between the algorithms and the data structures. We describe the principal abstractions comprising the GGCL, the algorithms and data structures that it provides, and we provide examples that demonstrate the use of GGCL to implement some common graph algorithms, as well as an implementation of the minimum degree sparse matrix ordering algorithm. Also discussed are particular implementation issues related to performance. Experimental results are presented to demonstrate that the novel lightweight implementation techniques and static polymorphism in GGCL results in code that is significantly more efficient than similar libraries written using the object-oriented paradigm. In the case of sparse matrix ordering, our implementation of the minimum degree algorithm has performance indistinguishable from that of a widely used (and highly tuned) Fortran implementation.

Speaker's web page: http://www.lsc.nd.edu/~lums/

Research web page: http://www.lsc.nd.edu/research/ggcl/

Institution web page: http://www.nd.edu/

November 12, 1999

Uniform Data Access Using GXD

Peter Vanderbilt
MRJ Technology Solutions

Email: pv@nas.nasa.gov

Abstract:

This talk will give an overview of GXD, a framework facilitating publication and use of data from diverse data sources. GXD defines an object-oriented data model designed to represent a wide range of things including data, its metadata, resources, and query results. GXD also defines a data transport language, a dialect of XML, for representing instances of the data model. GXD also provides a software library, prototyped in Java, that includes support for generating GXD-encoded entities and for interpreting these entities (potentially from many data sources) to create an illusion of a globally interconnected data space, one that is independent of data source location and implementation.

Institution web page: http://www.mrj.com/

November 2, 1999

Dyninst: Status and Future Evolution

Jeff Hollingsworth
University of Maryland, College Park

Email: hollings@cs.umd.edu

Abstract:

I describe a post compiler program manipulation called dyninstAPI which provides a C++ class library for machine independent program instrumentation during application execution. After a brief introduction to the topic, the talk will concentrate on presenting the new and recently proposed features in the API including support for type information, program modification, and use as an off-line binary editor. In addition, I will describe a new TCL-based command line tool we have developed to speed the creation of simple tools that use the API. I will also discuss the status of the DPCL interface being developed by IBM.

Speaker's web page: http://www.cs.umd.edu/~hollings/

Research web page: http://www.cs.umd.edu/projects/dyninstAPI/

Institution web page: http://www.umd.edu/

October 20, 1999

KULL System Architecture

Pat Miller

Scientific Computing Applications Division

Lawrence Livermore National Laboratory

Abstract:

The KULL project is a joint A-Division/X-Division effort to build a 3-D modeling code for simulating ICF capsules. The code is built on an object-oriented framework. Individual physics packages are written either in C++ or in Fortran. We make clever use of templatization and inheritance to promote code reuse and to maintain a very abstract view of the mesh geometries used by the physics packages. The entire code is "steered" from a Python interpreter with embedded physics objects. The presentation will include some architectural details of the project, including our use of parallelism, C++ features, and steering.

Institution web page: http://www.llnl.gov/

October 18, 1999

Diamond Eye: A System for Mining Large Image Collections

Michael Burl
Jet Propulsion Laboratory

Email: burl@aig.jpl.nasa.gov

Abstract:

Diamond Eye is a distributed software architecture that enables users (scientists) to analyze large image collections by interacting with one or more custom data mining servers via a Java applet interface. Each server is coupled with an object-oriented database and a computational engine such as a network of high-performance workstations. The database provides persistent storage and supports querying of the "mined" information. The computational engine provides parallel execution of expensive image processing, object recognition, and query-by-content operations.

Key benefits of the Diamond Eye architecture are: (1) the design promotes trial evaluation of advanced data mining and machine learning techniques by potential new users (all that is required is to point a web browser to the appropriate URL); (2) software infrastructure that is common across a range of science mining applications is factored out and reused; and (3) the system facilitates closer collaborations between algorithm developers and domain experts.

Speaker's web page:
http://www-aig.jpl.nasa.gov/public/mls/home/burl/home.html

Institution web page: http://www.jpl.nasa.gov/

October 15, 1999

Development of a Three-Dimensional Relativistic PIC Code for Studying the Production of Useful Electron Bunches

Roy Hemker
University of California, Los Angeles

Email: hemker@physics.ucla.edu

Abstract:

The advances in computational speed make it now possible to do full 3D PIC simulations of laser plasma and beam plasma interactions, but at the same time the increased complexity of these problems makes it necessary to apply modern approaches like object-oriented programming to the development of simulation codes. We report here on our progress in developing an object-oriented parallel 3D PIC code using Fortran 90. In its current state the code contains algorithms for 1D, 2D, and 3D simulations in cartesian coordinates and for 2D cylindrically symmetric geometry. For all of these algorithms the code allows for a moving simulation window and arbitrary domain decomposition for any number of dimensions. Recent 2D and 3D simulation results on the propagation of intense laser and electron beams through plasmas will be presented.

Institution web page: http://www.ucla.edu/

October 7, 1999

Semiclassical Quantum-Mechanical Modelling using Parallel PIC Methods

Dean Dauger
University of California, Los Angeles

Email: dauger@physics.ucla.edu

Abstract:

We have developed a multiparticle quantum-mechanical dynamic modeler based on a semiclassical approximation of Feynman path integrals. This approximation transforms the quantum problem into a problem of classical trajectories, thus enabling us to implement the simulation on parallel computers using mature Particle-In-Cell methods developed for plasma codes. With modern hardware, this should allow us to simulate multiple (100-1000) interacting quantum particles, treating the Hamiltonians semiclassically. Techniques that enable the numerical reliability of the simulation were developed to handle the grid-point nature of PIC methods and boundary conditions for the semiclassical case. One-dimensional single and multiparticle examples will be shown, with analyses of their energetic structure, and a technique for analyzing fermionic cases will be introduced.

Institution web page: http://www.ucla.edu/

October 5, 1999

Nonlinear Solver Challenges in Computational Fluid Dynamics

David Young
The Boeing Company

Email: david.p.young@boeing.com

Abstract:

Many important phenomena in fluid dynamics are highly nonlinear. Some examples include shock waves and boundary layer separation. There are well-known instances of non-uniquess and ill-posedness even for inviscid transonic flow. These difficulties are not ameliorated by solving the Navier–Stokes equations. Design optimization problems inherit all these challenges and add ill-posedness related to constraints. In this talk, we will discuss some of these difficulties and some computational strategies for dealing with them. These strategies include nonlinear elimination, nonlinear Schwartz methods, and the use of multi-point design. A rather complete description of our design optimization methodology will be given with a discussion of the implications of various methodological choices for global convergence to the solution of the problem.

Institution web page: http://www.boeing.com/

October 4, 1999

Nonlinear Solvers and Preconditioners in Groundwater Flow Simulations

Timothy Kelley
North Carolina State University

Email: Tim_Kelley@ncsu.edu

Abstract:

Newton–Krylov–Schwarz methods solve nonlinear equations by using Newton's method with a Schwarz domain decomposition preconditioned Krylov method to approximate the Newton step. In this talk we will discuss the design and implementation of Newton–Krylov–Schwarz solvers in the context of the implicit temporal integration on an unstructured three-dimensional spatial mesh of Richards' equation for flow in the unsaturated zone. The issues include nonsmooth nonlinearities, construction and efficient implementation of the coarse-mesh problem, and temporal integration. We conclude with a report of some preliminary numerical results.

Speaker's web page: http://www.math.ncsu.edu/~ctk/

Institution web page: http://www.ncsu.edu/

ISCR

Institute for Scientific Computing Research

ITS Lecture Series and Lecturer Biosketches

CASC Center for Applied Scientific Computing

August 30, 2000

The Design of Design

Frederick P. Brooks, Jr.
University of North Carolina

Email: brooks@cs.unc.edu

Abstract:

The processes of designing computers, programming languages, operating systems, and big applications seem to have a lot in common. More interestingly, these same commonalties are also found in the processes of designing buildings, and even organizations. In some fields, such as building architecture and industrial design, there is currently a lot of study of the design process. Perhaps by studying the process, across the various media in which designs are realized, we can learn how to better practice, teach, and design.

Engineers have a simplistic, deterministic model of their own design processes. As soon as one expresses the model, real designers acknowledge it and recognize its gross inadequacies. Architects rejected it decades ago. What is a better model?

Through the centuries, designers have debated the fundamental issue of whether design should be rational or empirical, a question with many corollaries and implications. This talk fiercely defends the empirical side.

Great works of art have typically come from one mind; many great works of design have the same property. Now, when time-to-market pressures and increased specialization of skills dictate design teams, how shall we achieve conceptual integrity in our designs? What are the implications for telecollaboration?

Finally, I shall argue that "Great designs come from great designers, not from great product processes," and that many product troubles come from terribly wrong consensus-based design processes.

Speaker's web page: http://www.cs.unc.edu/~brooks/

Institution web page: http://www.unc.edu/

Biographical Sketch:

Frederick P. Brooks, Jr

Frederick P. Brooks, Jr., was born in 1931 in Durham, North Carolina. He received an A.B. summa cum laude in physics from Duke University and a Ph.D. in computer science from Harvard, under Howard Aiken, the inventor of the early Harvard computers. Dr. Brooks then joined IBM, working in Poughkeepsie and Yorktown, New York, from 1956 to 1965. He was an architect of the Stretch and Harvest computers and then was the project manager for the development of IBM's System/360 family of computers and Operating System/360 software. For this work he received a National Medal of Technology jointly with Bob O. Evans and Erich Bloch.

In 1957 Dr. Brooks and Dura Sweeney patented an interrupt system for the IBM Stretch computer that introduced most features of today's interrupt systems. His System/360 team first achieved strict compatibility—upward and downward—in a computer family. Brooks coined the term "computer architecture" in relation to the System/360 family. His early concern for word processing led to his selection of the 8-bit byte and the lowercase alphabet for the System/360; he engineered many new 8-bit input/output devices and provided a character-string datatype in PL/I.

In 1964 he founded the Computer Science Department at the University of North Carolina at Chapel Hill and chaired it for 20 years. Currently, he is Kenan Professor of Computer Science at UNC. His principal research is in real-time, three-dimensional, computer graphics "virtual reality." His research has helped biochemists solve the structure of complex molecules and has enabled architects to "walk through" structures still being designed. He is pioneering the use of force display to supplement visual graphics.

Dr. Brooks distilled the successes and failures of the development of Operating System/360 in the book *The Mythical Man-Month: Essays in Software Engineering* (1975, 20th Anniversary Edition, 1995). He further examined software engineering in his well-known 1986 paper, "No Silver Bullet." In 1997, he and Professor Gerrit Blaauw published a major research monograph, *Computer Architecture: Concepts and Evolution*.

Brooks has served on the National Science Board and the Defense Science Board. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences. He has received the ACM A.M. Turing Award, the IEEE John von Neumann Medal; the IEEE Computer Society's McDowell and Computer Pioneer Awards, the ACM Allen Newell and Distinguished Service Awards, the AFIPS Harry Goode Award, an honorary Doctor of Technical Science from ETH-Zurich, and The National Medal of Technology (from President Reagan) in 1985.

June 9, 2000

Mathematics and Computers

Peter Lax

Courant Institute of the Mathematical Sciences, New York University

Email: lax@cims.nyu.edu

Abstract:

This talk will give an overview of the rise of computing in mathematics, and the infusion of mathematical ideas into computing in the past half century. The founding father of computational mathematics is von Neumann, who first set the agenda for the development of programmable electronic computers with large memory, as well as programming languages, schemes of discretising continuum equations, and algorithms for solving the discretised equations.

Everybody with the least familiarity with computing is aware of the continuing phenomenal improvement in the last fifty years of computer speed, memory capacity, software, and graphics capabilities. Problems that at one time have strained computing capacities could, a few years, later be done routinely, cheaply, and fast. But most people are unaware of how much this progress is due in equal measure to new mathematical ideas for tackling computational tasks. Striking examples include: shock capturing, multigrid, the fast Fourier transform, fast matrix multiplication, multiresolution, neural nets, etc. No doubt many more such mathematical shortcuts remain to be discovered.

von Neumann also realized that computing can do more than grind out by brute force numerical answers to concrete questions; it can discover new phenomena. Theorists can use computing in the manner of the experimentalists. In just such fashion did Fermi, Ulam, and Pasta discover recurrence in the dynamics of nonlinear chains, Kruskal and Zabusky solitons, Feigenbaum period doubling, Lorenz strange attractions, and Mandelbrot fractals. Perhaps the most astonishing discovery through numerical experimentation is Odlyzko's exploration of the connection between the distribution of the zeros of Riemann's zeta function and the distribution of the eigenvalues of random matrices. To prove rigorously what has been discovered by computation is a tremendous challenge of the future.

Computing has also made its mark as an aid in proofs. The most spectacular example is Haken and Appel's demonstration of the four-color theorem, where the computer is required to perform exact calculations. But computing is also an integral part of the logical structure of some analytical proofs, where the computations are only approximate. One can look forward to more and more computer assisted proofs. Have fun!

Institutional web page: http//www.nyu.edu/

Biographical Sketch:

Peter Lax

Peter Lax received his Ph.D. in 1949 from the Courant Institute of Mathematical Sciences at New York University. He then went to Los Alamos, where he had been stationed in the Army during the last year of World War II. He worked with John von Neumann, devising schemes for calculating flow with many interacting shocks. His first numerical experiments were carried out in 1952 on Los Alamos' MANIAC computer.

In 1954 he returned to the Courant Institute to join the Atomic Energy Commission Computing Center. He has spent the remainder of his professional career at NYU, making leading contributions to computational mathematics. His numerous honors include the National Medal of Science, the Lester R. Ford Award, the Chauvenet Prize, the Semmelweis Medal, the Wiener Prize, and the Wolf Prize. Many of his students have themselves become leaders in scientific computing.

Among his many contributions to the High Performance Computing community, one of the most important was his role in establishing the National Science Foundation (NSF) computing centers. He served on the National Science Board from 1980 to 1986 and chaired the "Lax Panel," which recommended that the NSF establish five national computing centers to provide university scientists access to the same level of high-performance computing available in the national laboratories. The articulation of this requirement in national science policy helped spur the birth of all subsequent HPCC programs, including ASCI.

May 12, 2000

How Shall We Program High Performance Computers?

Burton Smith
Tera Computer Company

Email: burton@tera.com

Abstract:

Uniprocessor computer architecture has traditionally been motivated by programming languages and operating systems, with benchmarks written in the usual languages also having some influence. In high performance computing the situation is curiously reversed, with architecture determining the principal characteristics of programming languages, operating systems, and benchmarks. The result has been chaos; a "software crisis" has been declared, and better tools for the development of parallel software have been demanded. The outlook for good tools is bleak without a new approach to the problem, which should include the engineering of computer systems with both system and application software in mind and the development of programming abstractions that are both effective and efficient on hardware we can build.

Speaker's web page: http://www.tera.com/company/management/smith.html

Research web page: http://www.tera.com/company/index.html

Biographical Sketch:

Burton Smith

Burton Smith is Chief Scientist of Cray Inc. He has held a variety of faculty, federal research, and corporate project leadership posts. He received the BSEE from the University of New Mexico in 1967 and the Sc.D. from MIT in 1972.

From 1985 to 1988 he was Fellow at the Supercomputing Research Center of the Institute for Defense Analyses in Maryland. Before that, he was Vice President of Research and Development at Denelcor, Inc., and was chief architect of the HEP, one of the earliest commercial parallel machines. He developed the Tera Multithreaded Architecture (MTA) and compiler, one of the most innovative parallel architectures ever, with 128 regularly scheduled threads per processor. The company he founded to commercialize the MTA, Tera Computer, bought Cray Inc. on April 27, 2000.

Dr. Smith is a Fellow of both the ACM and the IEEE, and he was winner of the IEEE-ACM Eckert-Mauchly award in 1991. His main interest is general-purpose parallel computer architecture.

He has visited Lawrence Livermore National Laboratory well over a dozen times, beginning in 1979.

April 12, 2000

Partly Random Graphs and Small World Networks

Gilbert Strang
Massachusetts Institute of Technology

Email: gs@math.mit.edu

Abstract:

It is almost true that any two people in the United States are connected by less than six steps from one friend to another. What are models for large graphs with such small diameters? Outstanding applications would be networks of neurons, or an electric power grid, or even (possibly) the world wide web.

Watts and Strogatz observed (in *Nature*, June 1998) that a few random edges in a graph could quickly reduce its diameter (longest distance between two nodes). We report on an analysis by Newman and Watts to estimate the diameter with an N-cycle and M random shortcuts, $1 << M << N$.

We also study a related model, which adds N edges around a second (but now random) cycle. The average distance between pairs becomes nearly $A \log N + B$. The eigenvalues of the adjacency matrix are surprisingly close to an arithmetic progression; for each cycle they would be cosines, the sum changes the spectrum completely.

We will discuss diameters and eigenvalues of the adjacency matrix for partly random graphs. We also report on the surprising eigenvalue distribution for trees (large and growing) found by Li He and Xiangwei Liu. And a nice work by Jon Kleinberg discusses when the short paths can be located efficiently by a decentralized algorithm—as on the web.

Speaker's web page: http://www-math.mit.edu/~gs/

Institution web page: http://www.mit.edu/

Biographical Sketch:

Gilbert Strang

Gilbert Strang was an undergraduate at MIT and a Rhodes Scholar at Balliol College, Oxford. His Ph.D. is from UCLA and upon graduation he has taught at MIT. He has been a Sloan Fellow and a Fairchild Scholar and is a Fellow of the American Academy of Arts and Sciences. He is a Professor of Mathematics at MIT and an Honorary Fellow of Balliol College.

Professor Strang has published a monograph with George Fix, "An Analysis of the Finite Element Method," and six textbooks:

- Introduction to Linear Algebra (1993,1998)
- Linear Algebra and Its Applications (1976,1980,1988)
- Introduction to Applied Mathematics (1986)
- Calculus (1991)
- Wavelets and Filter Banks, with Truong Nguyen (1996)
- Linear Algebra, Geodesy, and GPS, with Kai Borre (1997)

He served a two-year term as the President of the Society for Industrial and Applied Mathematics (SIAM), 1999-2000, during which tenure he helped lead SIAM in the formation of an Activity Group in Computational Science & Engineering and organized the first SIAM conference in CS&E.

ISCR

Institute for Scientific Computing Research

University Collaborative Research Program Subcontract Research Summaries

CASC Center for Applied Scientific Computing

Software Infrastructure for Multi-Tier Implementation of Structured Adaptive Mesh Hierarchies

Scott Baden
University of California, San Diego

Summary:

We have continued our investigation into scalable parallel programming methodology and performance tradeoffs arising in mixed-mode (message-passing, multi-threaded) programming of hierarchically constructed multi-tier computers. We have leveraged previous work with the KeLP system for manipulating hierarchical data structures, which is currently running on ASCI Blue Pacific. Our investigation targets structured adaptive mesh refinement, including hierarchical load balancing strategies. A uniformly refined multi-block code, including a load balancer, was developed first, followed by a structured adaptive mesh version. An existing implementation of an SAMR Poisson solver, written in a single-tier dialect of KeLP, was ported to hierarchical KeLP. The investigation will deliver a computational testbed permitting LLNL scientists to explore performance tradeoffs under mixed-mode parallelism for structured adaptive mesh refinement and a variety of other block-structured applications of interest to the laboratory. We are conducting comparative studies with the NPACI Tflops architecture, which is based on Power3 SMPs of an architecture similar to Livermore's ASCI White.

Principal LLNL contact: John May, CASC

High Resolution Parallelized Self-Organizing Maps

Jackson Beatty
University of California, Los Angeles

Summary:

This project was designed to create publicly available parallel algorithms for computing self-organizing neural networks. However, it also gave rise to two major new programs at UCLA. The first is a research program seeking to establish the computational basis of cortical information processing in the human brain. The second is a program of instruction in neuroinformatics at UCLA, funded by the Human Brain Project, which is the first such program in the United States.

Four new parallel algorithms for computing and extending Kohonen's self-organizing map (SOM) algorithm were created under this grant. The first was psom.fortran, a generalized version of the SOM algorithm coded in Fortran and utilizing the MPI library for parallel computation. The psom.fortran algorithm freed the Kohonen procedure from several of its arbitrary parameters and provided a general mechanism for understanding unsupervised competitive learning processes embedded within a two-dimensional surface, similar to that of the human cerebral cortex. The second algorithm was psom.c, which is algorithmically identical to psom.fortran but was completely rewritten because of MPI library difficulties with Fortran on Livermore Teracluster. The third algorithm we call Adaptasom, for adaptive coordinate visualization of self-organizing map formation. Adaptasom accepts information from either psom.c or psom.fortran as a map is being generated and displays it as movements of reference vectors in the self-organizing map space. Adaptasom is based on insights and algorithms developed by Rauber Andreas in his doctoral thesis at the Technical University of Vienna in 1996. The fourth algorithm is called Retina, which is a set of procedures for transforming calibrated digital photographs of natural images into a representation similar to that provided by the retina to the optic nerve and hence to human visual cortex. Retina provides the primary means of testing the Neuronal Empiricism Hypothesis and for discovering the common computational algorithm of the primate and human cerebral cortex.

For decades neuroscientists have studied the properties of cells in various specialized areas of the cerebral cortex and detailed their unique properties. Particularly well studied are the primary sensory areas, those cortical regions that receive the initial input from the different senses. Thus, the properties and organization of cells in the primary visual cortex are by now well known. Most of these cells respond best to lines of a particular orientation in a specific region of the retina. These cells form a part of a pinwheel, the segments of which contain cells favoring a specific particular organization. The pinwheels themselves are organized into a spatial mapping of the retina, so that each particular region of the cortex exclusively processes information from a unique portion of the retina.

Thus, it was of profound importance to learn that an identical pattern of neural organization emerges in auditory cortex when visual information is surgically redirected from visual to auditory cortex in newborn ferrets. The only possible explanation for this finding is that the cellular properties of visual cortex neurons are not genetically determined (as is usually assumed) but rather are extracted from information provided by the visual system using a computational algorithm that is common to all cerebral cortex. This is an extremely powerful idea, which forms the basis for what I term the Neuronal Empiricism Hypothesis. In its simplest

Summary (continued):

form, it simply states that the cortex learns to extract regular features from whatever sensory information it is provided, because such regularities provide information about real objects in the organism's environment.

The striking pattern of results that we are beginning to observe from our psom algorithm operating on the output of retina processing calibrated photographs of natural scenes suggests that the Kohonen algorithm (or some similar algorithm) may be capable of generating the observed properties and organization of cortical visual neurons simply by extracting regularities from the sensory stream provided by the retina. This finding, should it prove robust, would constitute a major advance in the field of computational neuroscience and provide deep evidence that there in fact exists a common algorithm for cortical computation.

This heavily computational but data-constrained approach to understanding the human nervous system also has educational implications. Our LLNL grant has provided us at UCLA with sufficient experience in large-scale computation to formulate a new program in neuroscience instruction at UCLA, called neuroinformatics.

One grand challenge facing science today is to transform neuroscientific data into facts and from these facts extract knowledge and understanding. Neuroinformatics offers a set of powerful tools by which this transformation of data may be accomplished. The new neuroinformatic tools include the creation of systematic, content-driven databases, sophisticated data visualization methods, unsupervised data-driven pattern extraction methods, neural modeling across a wide range of levels, and electronic collaboration. Similar methods have been applied with striking success to the (simpler) problems of gene and protein sequencing, an enterprise that created the discipline of bioinformatics. All of these new tools are examples of new information technology and—taken together—they have redefined the field of scientific computing.

Thus, neuroinformatics brings together the methods and concepts of neuroscience, informatics, and computation that are necessary to understand the human brain in all its complexity. Although computers have played important roles in neuroscience research for decades, it has only been in the last several years that the economics of computer manufacturing and marketing have made large-scale supercomputing technology available for the study of the brain and its functions. Neuroinformatics offers neuroscience new opportunities for extracting understanding from data and meaning from measurements.

With the support of the Human Brain Project, we are beginning to prepare a new generation of informatically oriented neuroscientists by specifically introducing them to neuroinformatic methods and principles while they are still undergraduates, and to make the curricula by which we teach neuroinformatics to our students freely available to colleges and universities throughout the world. This instructional program grew directly from the computational studies that we carried out over the last two years at LLNL.

Principal LLNL contact: Farid Dowla, Electronics Engineering Technology

Semiclassical Quantum-Mechanical Modeling using Parallel PIC Methods

John Dawson and Viktor Decyk

University of California, Los Angeles

Summary:

We have been successful in modeling many particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas. Issues regarding the derivation and implementation of these techniques are described briefly, a method to address quantum particle statistics is introduced, and studies of two quantum-mechanical problems—a two-electron, one-dimensional atom and electrostatic quasi-modes due to quantum effects in a hot electron plasma—are discussed.

We have continued the collaborative program with Denis Hewett at LLNL on the use of parallel computers for quantum physics. The effort concentrates on building models and developing techniques to simulate systems of practical interest. Our focus is on multiparticle systems where quantum effects are significant.

UCLA has pioneered particle methods (PIC) for modeling plasmas. For over six years, we have developed efficient and effective methods for using parallel computers for carrying out PIC simulations (90% parallel efficiency, 40% of estimated peak speed). At present we are able to handle over 200 million particles in three dimensions using self-consistent full electromagnetic fields. Our methods are robust and portable and have run successfully on a wide range of computers (e.g., Cray-90s, T3Ds, T3Es, and SP2s).

The focus of the project has been to apply PIC techniques to the semiclassical approach and develop a dynamical model of many particle quantum systems. This model would allow detailed investigation of optical properties, ionization potential, conductance, and a host of other experimentally determined material properties; it has predictive capabilities useful for the design and understanding of devices where quantum effects are important. The results have been very promising. This work is the part of the Ph. D. thesis of Dean E. Dauger.

We have made good progress towards a dynamic quantum mechanical modeler and have accomplished two major milestones: we have successfully transformed a plasma code into a quantum code, and we have applied this code to a range of quantum problems, including multiple particles.

The theoretical basis for the methods developed here is the Feynman path integral, which can predict the time evolution of a wavefunction using a sum over all possible paths. However, based on Van Vleck's Green's function, the path integral may be reduced to a sum over all possible classical paths, using the semiclassical approximation. This approach is easily extended to multiple dimensions and to include full electromagnetism. It turns out the primary cost of this method is in tracing the classical trajectories through electromagnetic fields.

Summary (continued):

We considered how to handle interactions between multiple particles. For typical interactions described in a multiparticle Hamiltonian, we know how to derive an expansion for the effective Hamiltonian for any one particle. The ordering coefficient of this expansion is proportional to the time of the wavefunction push, implying a recommended limit for how long to evolve the wavefunction before reconstructing the wavefunction and updating the fields. Spin-spin interactions can be incorporated very easily into this structure.

Our investigation has also focused on the accuracy of the techniques essential to this project. We have produced and examined a precise derivation of the semiclassical method as it applies to our problem and have addressed a significant issue in using semiclassical methods. When time evolving one quantum particle over one quantum time step, the derivation describes that the sum over all paths between an initial and a final point simplifies to one classical path. This derivation also produces a coefficient, which leads to caustics and related potential difficulties in using semiclassical methods. Classical path times in previous use of semiclassical methods were long enough for caustics to pose a significant challenge. However, because our approach repeats many quantum time steps to evolve the wavefunctions, path times in our use of these methods are much shorter, making caustics and related concerns unlikely in the systems with which we are interested.

Another issue we investigated was the application of these techniques to systems exhibiting fermion statistics. This issue is of importance when applying these techniques to the simulation of electrons.

The first system we applied this technique to was two fermions in an infinite square well. We chose this potential system because it had a well-known single particle energy spectrum yet did not have the regularity of the simple harmonic oscillator system. Peaks in the energy spectra were found where we thought they should be, at energies that correspond to two-particle states of differing quantum number. More importantly, no peaks were found where we thought they should not be, that is, at those energies that would contradict the Pauli exclusion principle and energies in between expected energies.

We investigated further fermionic multiparticle problems. The next system was a one-dimensional atom: two mutually repulsive fermionic electrons attracted to a central, positive, one-dimensional charge. This system was chosen for its appropriateness to the code while being significantly more difficult to solve analytically. As with the infinite square well case, we found the energy spectra in the one-electron case and the two-electron case, and we found the ground-state energy to be comparable to what we could predict analytically using variational methods.

Summary (continued):

There is a great deal of what we can do with the code, which we have used for problems with over a hundred quantum particles. We intend to pursue further analysis of the quantum fluctuations problem, perhaps extending it to predict solar reaction rates. Also, it is possible to investigate particular issues regarding the implementation. One issue is to resolve why the boundary conditions for the infinite square well require the virtual classical particles bounce one half grid beyond the boundary. Another is to consider alternative decompositions for better parallelism, as the current decomposition over space is not ideal when the number of grids is low. Third, the current virtual particle sampling scheme blankets phase space. There may be ways of simplifying or optimizing this distribution to reduce the computational cost spent to evolve a wavefunction.

Further in the future, a number of additional features should offer a more complete modeling of quantum mechanics. Other higher order effects can be included in the code. Techniques at implementing two and three dimensions and full electromagnetism should carry over well from the plasma PIC codes. Along with higher dimensions, we can fully incorporate spin phenomena. With these additions, the code would be able to model multielectron atoms, chemical reactions, quantum electronics, and solid state physics, among a host of addressable physical problems.

We are applying the code to a problem concerning fluctuations in the electrostatic spectrum in a hot electron plasma due to quantum mechanics. A particular analysis of this problem may have important consequences for astrophysical predictions about stellar evolution and the early universe. It has been proposed that the energy density in the electromagnetic spectrum of a hot plasma when including quantum effects is significantly different than what is predicted using classical theory alone, which we intend to test.

Principal LLNL contact: Dennis Hewett, CASC

Development of Bioinformatic Tools for the Integration and Analysis of Microarray Data

Jeffrey Gregg
University of California, Davis

Summary:

DNA microarrays offer the ability to determine gene expression values of hundreds to thousands of genes simultaneously, offering insight into the complex behavior of the cells being studied. However, there currently exists no robust system for cataloguing and analyzing the massive amounts of data being generated in these micorarray studies. Our principal objective has been to develop, evaluate, and implement a flexible database schema and develop large-scale computational methodologies for the analyses of expression microarray data generated by the projects already initiated in Dr. Gregg's and Dr. Wyrobek's research laboratories at UC Davis and LLNL, respectively. To this end, we are designing and implementing appropriate database structures for support of gene expression data at UC Davis and LLNL and we are developing analytical tools for the interpretation of microarray data and its correlation with cellular biomarker measurements and clinical outcomes. Development of an effective database management system and analytical tools will allow interpretation of biologically relevant gene expression data that will impact the understanding of human disease.

Principal LLNL contact: Andrew Wyrobek, Biology and Biotechnology

Tailor-made Virtual Libraries

B. S. Manjunath
University of California, Santa Barbara

Summary:

Different people can differ significantly on the similarity of two multimedia objects, such as still images or videos. To an urban planner, two regions are similar because of the distribution of buildings, parks, and roads. However, when an analyst in a forest department looks at these images of urban regions adjacent to forests, her discrimination could be based on a different set of attributes, including factors such as a possible fire hazard because of dead wood content. To a botanist two pictures can be similar because they contain views of the same flower, while to a casual person the images might be dissimilar because the backgrounds are quite different. It is difficult, in general, to extract image and video features that support all possible user perspectives. This in turn makes it hard to access such multimedia databases by content.

We have outlined a strategy that could lead to the efficient access of large multimedia databases that is tailored to individual perceptual criteria. Our method is capable of learning the user's perceptual metric and using it to access the database by content. By learning these personal views and making them available in a useful manner, the knowledge and information content of the database grows. The database is no longer just a repository of bits and bytes, but also of knowledge and information that grows dynamically and is available in a concrete useful manner. Over time a layer of meta-data arises over the database providing multiple views of what was once a monolithic object. Thus, our objective is to develop a strategy for efficient access of large multimedia databases that is tailored to individual perceptual criteria.

The focus of this research matches quite well with the ISCR focus area of large-scale data mining and pattern recognition. We collaborate with Imola Fodor, Center for Applied Scientific Computing at LLNL, on this research. Dr. Fodor's group is developing scalable algorithms for the interactive exploration of large, complex, multi-dimensional scientific data, and the UCSB group's work on integrating dimensionality scaling and learning perceptual metric directly contributes to this effort. The techniques developed in this project have a wide range of applications. To compute the mapping from the image features to the perceptual space, new tools for dimensionality reduction have been developed. The mapping from the user's perceptual criteria to a low dimensional vector space is useful in developing fast indexing and retrieval strategies in large scientific databases, and in data mining.

We are building a software prototype for large scientific image databases (such as satellite or aerial imagery). The prototype will be used to test various techniques for the three components, i.e., creation of feature vectors (XVs), computing the c-Maps, and adaptive indexing.

Principal LLNL contact: Imola Fodor, CASC

Development of a Three-Dimensional Relativistic Particle-in-Cell Code for Studying the Production of Useful Electron Bunches Using Ultra-Intense Laser Pulses

Warren B. Mori
University of California, Los Angeles

Summary:

We have developed a parallelized three-dimensional relativistic particle-in-cell (PIC) code for studying the production of useful electron bunches using ultra-intense laser pulses. The electron bunches are produced when the radiation pressure of an intense laser either directly accelerates electrons or excites a plasma wave wake which accelerates electrons. Ultra-short electron bunches could be useful for injection into high-energy particle accelerators, for radiation sources, and for the fast ignitor fusion concept. A reliable and robust code will be invaluable towards the realization of such ultra-short electron bunches in a laboratory.

Our object-oriented Fortran90 parallelized PIC code has been used to model intense electron and laser pulses propagating through tenuous plasmas. Based on large-scale runs, we have carried out performance optimizations. We plan to add additional current/charge deposition schemes and to continue to work with Dr. Langdon's group to benchmark various 3D parallel algorithms.

Principal LLNL contact: Bruce Langdon, X Division

Sensitivity and Uncertainty Analysis for Large-Scale Differential-Algebraic Systems

Linda Petzold
University of California, Santa Barbara

Summary:

Sensitivity analysis of large-scale differential-algebraic (DAE) systems is important in many engineering and scientific applications. Sensitivity analysis generates essential information for design optimization, parameter estimation, optimal control, model reduction, process sensitivity and experimental design. Our recent work on methods and software for sensitivity analysis of DAE systems has demonstrated that forward sensitivities can be computed reliably and efficiently via automatic differentiation in combination with DAE solution techniques designed to exploit the structure of the sensitivity system. The DASPK3.0 software package was developed for forward sensitivity analysis of DAE systems with index up to two, and it has been used in sensitivity analysis and design optimization of several large-scale engineering problems.

Some problems require the sensitivities with respect to a large number of parameters. For these problems, particularly if the number of state variables is also large, the forward sensitivity approach is intractable. In the past year we have begun an investigation of adjoint (reverse) sensitivity methods, which have the potential to overcome this limitation. We have derived the adjoint system for general DAEs and investigated some of its fundamental properties. In particular, we have derived the adjoint DAE systems by a variational method, along with conditions for the consistent initialization of index-1 and Hessenberg index-2 adjoint DAEs. We have investigated the stability of the adjoint DAE system and shown that for semi-explicit DAE systems the adjoint DAE is stable whenever the original DAE is stable. For fully implicit DAE systems this is no longer true, but we can show that the stability is preserved by an augmented form of the adjoint system. Thus, it is this system that we discretize for the numerical solution.

Work has begun and is now well underway on an adjoint DAE solver, DASPK-ADJOINT, which uses a modified DASPK3.0 to solve the forward and reverse problems. Some preliminary work has been done on characterizing the class of well-posed sensitivity problems for DAEs, and on analysis of the numerical stability for fixed stepsizes of the BDF methods applied to the augmented adjoint DAE. Complexity analysis and preliminary numerical results indicate that the adjoint sensitivity method is advantageous over the forward sensitivity method for applications with a large number of sensitivity parameters and few derived functions.

During the course of this work, the PI visited LLNL several times and gave a lecture at the LLNL CASC Nonlinear Solvers Workshop in July 2000.

Principal LLNL contact: Steven Lee, CASC

Global Simulation of the Earth's Magnetosphere with Adaptive Mesh Refinement

Joachim Raeder
University of California, Los Angeles

Summary:

We have added structured adaptive mesh refinement capabilities to the UCLA Global Geospace Circulation Model (UCLA-CGCM) in collaboration with the SAMRAI project team in the Center for Applied Scientific at LLNL. The UCLA–GGCM is a global circulation model of Earth's magnetosphere and ionosphere that has been used for several years to study the interactions of the solar wind, magnetosphere, and ionosphere. The UCLA–GQCM with local adaptive mesh refinement supports a grid resolution in critical regions that is about two orders of magnitude better than the base code. This additional resolution allows us to address problems that were previously elusive to global modeling, such as flux transfer events, the formation of boundary layers, and the formation of thin current sheets in the late substorm growth phase.

SAMRAI is a general software support framework for structured adaptive mesh refinement (AMR) applications on parallel high performance computing hardware. In collaboration with the SAMRAI group, we will continue to implement a new GGCM code that integrates the existing numerical kernels from our current application into the SAMRAI framework. We will in particular address the multiphysics integration of the ionosphere submodel, and issues of visualization and the efficiency of the time integration algorithm in the adaptive mesh context.

Principal LLNL contacts: Richard Hornung and Scott Kohn, CASC

Theoretical Foundations of Single Bubble Sonoluminescence

Andrew Szeri
University of California at Berkeley

Summary:

Sonoluminescence (SL) is the production of light from micro-bubbles driven into extreme volume oscillations by a strong acoustic field. The temperatures and pressures in current experimental efforts are intense enough to produce, momentarily within the bubble, a cool dense plasma which emits light like a tiny star from the surface of an opaque core. The size of the light emitting region, and hence the luminosity of the bubble, depend sensitively on the physical characteristics of the material inside the bubble, and on the forcing to which it is subjected. Significant progress has been made on the theoretical understanding of SL, including the consequences of material selection, long-term chemical effects on the physics, and modeling techniques. We developed an accurate but fast computational model of heat and mass transfer in SL bubbles, based on a more detailed but very time-consuming model. The reduced model was subsequently applied to the study of different choices of materials (cryogenic noble gas systems, exotic low-vapor-pressure liquids) and the long-term consequences of chemical activity (the dissociation hypothesis).

Just prior to the project, we had written a detailed simulation to elucidate the role of water vapor in SL. We found that significant amounts of water vapor enter the bubble during the long slow expansion, as other researchers have shown, but that much water vapor is incapable of diffusing to the bubble wall in time to exit the bubble at the collapse. At higher and higher amplitudes of the driving, more and more water vapor is trapped in the bubble at the collapse. This water vapor is involved in chemical reactions, which are accelerated by the compression heating and high density of the bubble. Because the reactions are endothermic, they reduce the heating of the bubble, as does the lowering of the ratio of specific heats of the gas mixture. This is strongly coupled to observables: the number of photons and pulse width of light emitted. That study gave us the foundation we needed to develop a much faster model of reduced complexity. In this model, we take account of heat and mass transfer, as well as chemical reactions, but the fields of interest (pressure, temperature, and composition) have been averaged across the volume of the bubble contents. This leads to ordinary differential rather than partial differential equations. Despite the *simplicity* of the reduced model, comparison of results with the detailed model showed excellent agreement in terms of temperatures achieved, OH radical production, water vapor trapped, etc.

Summary (continued):

During the course of the research, in collaboration with William Moss of LLNL, we realized that the usual corrections for "compressibility" to the Rayleigh–Plesset equation for bubble dynamics had missed an important first-order effect. The usual corrections account for *liquid* compressibility, i.e., sound waves in the liquid, which contribute to the damping of the bubble dynamics by the radiation of acoustic energy. We realized that of even greater importance was compressibility of the gas, i.e., development of a nonuniform pressure field during rapid collapse or expansion of the bubble. We developed a simple correction to account for this effect, which yields much better agreement with experimental results for strongly forced (SL) bubbles.

We have also made use of the reduced model in the study of cryogenic systems. There was speculation that, for example, a Helium bubble in liquid Argon might be an ideal system for SL. The reason is that there would be no reduction of the internal temperature due to thermodynamic or chemical processes. Unfortunately we found that such bubbles would likely be overwhelmed by the large quantity of vapor that flows in during the expansion. We studied in addition other non-aqueous systems. These are challenging computations, as a multi-dimensional parameter space must be investigated in order to determine the narrow range of parameters where SL can be achieved.

As a final application of the reduced model we are studying the dissociation hypothesis of Lohse and coworkers. The idea is that diatomic components of air in an SL bubble undergo chemical reactions, resulting in soluble products that leave the bubble nearly pure Argon. This happens over several seconds, it is thought, which corresponds to perhaps 50–60,000 complete bubble oscillations. This places extreme demands on our computational model. We have developed the first simulation capable of computing this effect, but we must speed up the calculation further to develop data for comparison with experiments over a range of parameters.

Principal LLNL contact: William C. Moss, Geophysics and Global Security

ISCR

Institute for Scientific Computing Research

ISCR Subcontract Research Summaries

CASC Center for Applied Scientific Computing

Scalable Nonlinear Iterative Methods for Partial Differential Equations

Xiao-Chuan Cai
University of Colorado, Boulder

Summary:

We conducted a six-month investigation of the design, analysis, and software implementation of a class of singularity-insensitive, scalable, parallel nonlinear iterative methods for the numerical solution of nonlinear partial differential equations. The solutions of nonlinear PDEs are often nonsmooth and have local singularities, such as sharp fronts. Traditional nonlinear iterative methods, such as Newton-like methods, are capable of reducing the global smooth nonlinearities at a nearly quadratic convergence rate but may become very slow once the local singularities appear somewhere in the computational domain. Even with global strategies such as line search or trust region, methods for $F(u)=0$ often stagnate at local minima of $||F||$, especially for problems with unbalanced nonlinearities, because the methods do not have any built-in machinery to deal with the unbalanced nonlinearities. To find the same solution of u^* of $F(u)=0$, we solve, instead, an equivalent nonlinearly preconditioned system $G(F(u^*))=0$, whose nonlinearities are more balanced. In this project, we proposed and studied a parallel nonlinear additive Schwarz-based nonlinear preconditioner and showed numerically that the new method converges well even for some difficult problems, such as high Reynolds number flows, where a traditional inexact Newton method fails by stagnation.

We developed a general theoretical framework of nonlinearly preconditioned inexact Newton methods, which we presented at the CASC-sponsored Workshop for Scalable Nonlinear Solvers in July 2000 and which we have submitted for publication in *SIAM J. Scientific Computing*. As an example of the general framework, we formulated and tested in parallel on the Compass Alpha Cluster a nonlinear additive Schwarz preconditioner, called ASPIN (for Additive Schwarz Preconditioned Inexact Newton). Our test results for a viscous incompressible flow problem showed that the method is fast and insensitive to the high Reynolds numbers, where boundary layers appear. The method has great potential for other nonlinearly difficult problems, such as (1) problems whose solutions have local singularities such as shocks or nonsmooth fronts; and (2) multi-physics problems with drastically different stiffness that require different nonlinear solvers based on a single physics submodel. Although the method offers scalable nonlinear convergence in the limit of "bad parameters" such as Reynolds numbers, as with any Newton method, it needs to be coupled on the inside to a scalable linear solver. There exist many fruitful possibilities for combining ASPIN with the parallel linear solver software being developed at LLNL.

Developing a Tuned Version of ScaLAPACK's Linear Equation Solver

Jack Dongarra
University of Tennessee

Summary:

The LINPACK Benchmark has been used as a yardstick in measuring the performance of the Top 500 installed high-end computers. This benchmark was chosen because it is widely used and performance numbers are available for almost all relevant systems. The LINPACK Benchmark solves a dense system of linear equations. For the Top 500, the benchmark allows the user to scale the size of the problem and to optimize the software in order to achieve the best performance for a given machine. This evaluation does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations. Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good check of peak performance of a system.

By measuring the actual performance for different problem sizes *n*, a user can get not only the maximal achieved performance R_{max} for the problem size N_{max}, but also the problem size $N_{1/2}$ where half of the performance R_{max} is achieved. These numbers, together with the theoretical peak performance R_{peak}, are the numbers given in the Top 500 ranking. In an attempt to obtain uniformity across all computers in performance reporting, the algorithm used in solving the system of equations must confirm to the standard operation count for LU factorization with partial pivoting. In particular, the operation count for the algorithm must be $2/3\ n^3 + O(n^2)$ floating point operations.

As part of this project, we developed a version of the benchmark based on the hardware of the ASCI Blue Pacific system to achieve high performance. It was our goal to produce the fastest implementation that will take advantage of the hardware and software infrastructure on the Blue Pacific machine in achieving this mark. In order to accomplish this goal, our team drew on the expertise in this field developed with our work on the ScaLAPACK and ATLAS projects. We collaborated with researchers at LLNL to produce software for the Linpack benchmark that is fully optimized for the ASCI-Blue Pacific system.

The second part of this project, ATLAS, involves using timings coupled with code generation to automatically adapt linear algebra computations to run optimally on varying architectures. The expertise developed in this project will be leveraged in order to make machine-specific performance tweaks for the computations involved in the benchmark, in particular, tunings for various levels of caches, floating point register and unit usage, and optimal use of threading. We presented a tutorial at LLNL on the ATLAS optimization implementation and described ways to optimize numerical software in general.

Adaptive Mesh and Algorithm Refinement with Discrete Simulation Monte Carlo

Alejandro Garcia
San Jose State University

Summary:

In recent years there has been considerable interest in the physics of mixing a shock-accelerated interface between two gases. This simplified problem is of considerable practical importance for the design of inertial confinement fusion (ICF). Undesirable mixing occurs between the ablator and the fuel as a result of the Richtmyer–Meshkov instability, which occurs when the shock passes through the interface between the two gases. This mixing is detrimental to the efficiency of thermonuclear fusion, potentially reducing the yield below the break-even point.

Due to the inherent nonlinear nature of the problem, numerical simulation is one of the most important tools available to us. However, the broad range of scales, from nanometer (shock thickness) to ICF pellet size, causes a significant number of difficulties both from the algorithmic and the modeling side. The full Richtmyer–Meshkov problem is not only a multiscale problem but is also an example of a multiphysics problem in which different physics is required at different scales (e.g., non-hydrodynamic transport at molecular scales).

This project has developed a new technique for studying the Richtmyer-Meshkov problem, specifically designed for multi-physics: Adaptive Mesh and Algorithm Refinement (AMAR). The AMAR framework extends conventional mesh refinement to molecular scales by introducing a particle algorithm at those scales. Specifically, for gas dynamics, direct simulation Monte Carlo (DSMC) is the preferred simulation technique used for the study of shocks. Being a particle-based algorithm, DSMC is meshless, which is an important advantage to a problem that requires the tracking of the interface between two materials; in continuum simulations the algorithmic requirements for the tracking of the interface introduce limitations on the problems studied, or worse, inaccuracies. DSMC captures all the physics including the effects of molecular diffusion on the interface, the structure of the shock profile, and the dynamics of the shock interface interaction.

Our already demonstrated AMAR methodology for single species gas flows was extended to simulate multiple species. Our DSMC routines have been coupled to an Euler solver designed by the SAMRAI team at LLNL and the adaptive mesh refinement implemented using the SAMRAI framework. The new AMAR code will be validated by comparison with earlier AMAR studies and with DSMC results for the Richtmyer–Meshkov instability obtained jointly with Prof. N. Hadjiconstantinou of MIT.

Research on Parallel Adaptive Finite Element Methods

Michael Holst
University of California, San Diego

Summary:

In this project we studied several fundamental issues arising in the parallel adaptive solution of linear and nonlinear elliptic and parabolic PDEs using multilevel algorithms. We focused our efforts on a new approach described in the paper "A New Paradigm for Parallel Adaptive Mesh Refinement" by Bank and Holst. This new approach requires almost no communication to solve an elliptic equation in parallel, and therefore it has the potential to scale much more efficiently on massively parallel computers than do more traditional algorithms. The algorithm has an inherently multilevel structure, in that a sequence of problems on a refinement hierarchy of meshes is solved during the course of the calculation. In particular, the algorithm has three main components: (1) We solve a small problem on a coarse mesh and use a posteriori error estimate to partition the mesh. (2) Each processor provides the complete coarse mesh and instructs to solve the entire problem, but its adaptive refinement is largely limited to its own assigned mesh partition. (3) A final mesh is computed using the union of the refined partitions provided by each processor. The mesh is regularized into a global conformal mesh, and a final solution is computed using a standard overlapping domain decomposition method or a parallel multigrid method.

In certain circumstances the third step can be avoided, leading to an algorithm with no communication beyond that required to synchronize the processors at the beginning of the calculation. Unfortunately, to avoid this third step and still preserve the approximation quality of the resulting solution, a substantial amount of subdomain overlap must be enforced using heuristic techniques. One of our goals for this project was to study algorithms for performing the third step of the Bank and Holst algorithm efficiently on massively parallel computers with the use of parallel algebraic multilevel and multigraph methods, and to develop techniques for skipping the third step, which would avoid the need to use heuristic overlapping techniques. One breakthrough was the realization that the use of dual-problem error indicators provided an automatic overlapping technique that: (1) involved no heuristics, and (2) is optimal in a certain sense (the amount of required overlap is minimized). Portions of this work were presented at the *Root Finder's Ball* held in Pleasanton, CA, during our visit to LLNL, and portions were also presented in a CASC seminar during the visit.

Automated Diagnosis of Large Scale Parallel Applications

Karen L. Karavanic
Portland State University

Summary:

We are studying the application of experiment management techniques to the problems associated with gathering, storing, and using performance data; our goal is achieving completely automated diagnosis of application and system bottlenecks. In this project we focused on incorporating heterogeneous data from a variety of tools, applications, and platforms, and on designing novel techniques for automated performance diagnosis.

The first task includes representation and API for use and storage of performance data from a variety of tools, applications, and collection methods (tracing or profiling, online or offline, etc.). Eventually we will accommodate data gathered and used in both correctness and performance debugging. We investigated several currently available parallel performance tools: Vampir, Jumpshot, Paradyn, Pablo, AIMS, VT, Scala, and Guideview. We implemented extensions for our own PPerfDB prototype to incorporate resource descriptions and performance data gathered with VT, Vampir, and Jumpshot on the LLNL ASCI Blue and Compass platforms and a locally available cluster of SPARC Solaris workstations. This was the first inclusion of data from trace-based tools into the experiment management prototype.

The second task is developing new techniques for the diagnosis itself. Project member Reena John investigated the Paradyn performance consultant in the context of historical data and heterogeneous resource hierarchies. We began designing a diagnostic framework that will incorporate automatically determined thresholds. This expands upon earlier work in which we developed a new automated diagnosis technique to incorporate prior knowledge.

The Experiment Management paradigm is a useful approach for designing a tool that will automatically diagnose performance problems in large-scale parallel applications. The ability to gather, store, and use performance data gathered over time from different executions and using different collection tools enables more sophisticated approaches to performance diagnosis and to performance evaluation more generally. We will continue our efforts by further development and analysis of online diagnosis using historical data, and by investigating performance data and diagnosis gathered from mixed MPI/OpenMP applications.

Discretization and Iterative Solution Techniques for Elliptic Problems on Non-Matching Grids

Raytcho Lazarov and Joseph Pasciak

Texas A&M University

Summary:

The construction, analysis, and numerical testing of efficient solution techniques for solving elliptic PDEs that allow for parallel implementation are the focus of our research. We have investigated and analyzed a number of discretization and solution methods for solving second-order elliptic problems that include mortar and penalty approximations and domain decomposition methods for finite elements and finite volumes. Techniques for parallel domain decomposition algorithms in the framework of PETSc and HYPRE have been studied and tested. We have implemented hierarchical parallel grid refinement and adaptive solution methods and have tested them on various model problems.

Discretization methods for PDEs on non-matching grids are important in applied scientific computing. They provide great flexibility in the grid generation process, increase the portability of various approximation methods and computer implementations, enhance coarsening strategies in parallel algebraic multigrid methods, and provide a natural and practical means of *a posteriori* error analysis in parallel domain decomposition methods. Our research focuses on the analysis and testing of some methods from this promising approach in order to make competitive implementations in engineering codes for complex applications.

We have completed a comprehensive study and testing of various mortar spaces that can be used in non-conforming domain decomposition methods. This implies that meshes for individual subdomains do not necessarily match along the interface, and the continuity of the solution along the subdomain interfaces is imposed in a weak sense. The mortar space provides the "glue" for this process. However, most of the finite element approximations of the mortar space used in the mortar finite element method have been related to the traces of the finite element spaces on the interfaces, which result in continuous functions. We have constructed mortar spaces by using the dual bases approach and some finite volume ideas. Such constructions are relatively simple, valid for general meshes, and lead to local computations. We have constructed discontinuous local dual basis functions for the mortar space that led to diagonal mass matrices and, hence, simple computational schemes. Our mortar spaces have been tested on a series of three-dimensional problems that include second-order elliptic equations and problems of linear elasticity.

The interior penalty method aims at eliminating the need for additional (Lagrange multiplier or mortar) spaces and imposes (approximately) the required continuity across the interfaces by appropriate penalty terms. In our approach the jumps in the values of the functions along these interfaces is penalized in the variational formulation. For smooth solutions we lose the optimal accuracy due to lower approximation at the interface, but on the other hand we produce symmetric and positive definite discrete problems which have optimal condition numbers.

Summary (continued)

We also addressed the issue of constructing preconditioners for composite non-matching grid discretizations. We proposed and investigated an interface domain decomposition preconditioner, which is spectrally equivalent to the reduced interface algebraic problem. We have tested both the accuracy and the preconditioned method on a series of model problems. A similar procedure for combining mixed finite element problems does not require a mortar space. Methods for solving the corresponding composite problems have been proposed, studied, and tested.

We also worked on the problem of multilevel grid refinement and error control for both finite volume approximations and penalty domain decomposition methods. This work directly connects with research in CASC on developing and testing of parallel algorithms (J. Jones, V. Henson. R. Falgout, U. Meier, C. Tong, and P. Vassilevski). We have developed two- and three-dimensional codes for parallel adaptive grid refinement that produce nested (and matching) grids. The resulting multi-level structure is used to define multigrid preconditioners. In collaboration with P. Vassilevski and C. Tong the software was incorporated into the HYPRE preconditioner library. We use the Finite Element Interface (FEI) specification, which provides a layered abstraction that minimizes the concern for the internal details in the HYPRE library. The resulting software was used for testing various ideas and strategies in the a posteriori error analysis and error control for convection-diffusion-reaction problems in 3D domains with complex structure.

We shall now focus on developing multigrid-like (MG, AMG, multilevel) algorithms that exploit coarse finite element spaces that do not necessarily match across subdomain boundaries. In the first algorithm, which is based on geometric finite element discretization, the refinement, local and adaptive, will be the main direction. In order to allow for efficient parallelization, we will investigate approaches of independent refinement in each subdomain, thus facing the problem of non-matching grids across the subdomain interfaces. In the dual approach, which is related to AMGe (algebraic multigrid method for finite element problems), we need to create coarse spaces (or coarse grids) in parallel, and one possibility is to perform any available sequential coarsening in each subdomain, thus creating non-matching coarse grids across the subdomain boundaries. In either approach we have to resolve the issue of matching the independently created problems at the intermediate grids.

In summary, we are developing efficient parallelizable, scalable algorithms for treating either the geometrically refined spaces or algebraically coarsened grids. The research offers alternatives to the existing efforts in CASC for parallel AMG(e) method, which might be more efficient since it exploits additional properties of the matrices.

The AMR++ Library

Bobby Philip

Front Range Scientific Computations, Inc.

Summary:

Our project focused on the following objectives:

- Development and testing of the AMR++ adaptive mesh refinement library to ensure robustness.
- Development and implementation of multilevel elliptic solvers (FAC, AFAC, AFACx) on adaptively refined Cartesian grids using AMR++.
- Initial extension of these algorithms to overlapping logically rectangular curvilinear grids.
- Proof of the level independence of the condition number of the AFACx algorithm.
- Documentation of the C++ code.

Adaptive mesh refinement (AMR) computations are complicated by their dynamic nature. The development of elliptic solvers for realistic applications is complicated by both the complexity of the AMR and the geometry of the problem domains. The AMR++ library represents an attempt to provide users with the framework to develop their own AMR applications from existing single grid solvers. Ensuring that AMR++ was a robust piece of software, and extending the interface to handle AMR on overlapping grids such as those present within the Overture framework, were principal objectives of this work. To this end, extensive development and testing were performed. The code was also documented in detail. The FAC, AFAC, and AFACx algorithms were implemented on Cartesian grids and the possibility of extending these algorithms to overlapping grids was explored. Initial research in this direction is being currently pursued. Proof that the condition number of the AFACx operator is independent of the number of refinement levels was also developed.

Scalable Solvers and Applications in Advanced Materials

Calvin J. Ribbens

Virginia Polytechnic Institute and State University

Summary:

The purpose of the scalable linear solvers and *hypre* projects is to develop scalable algorithms and software for the solution of large, sparse linear systems of equations using massively parallel computers. In order to meet the needs of a variety of current and future users, an important goal of the *hypre* package and framework is interoperability and extensibility. Toward that end, the PI worked with CASC's Andrew Cleary and Jeff Painter to formalize and extend an object model for *hypre*. From the beginning, *hypre* has been designed in an object-oriented style. The existing library includes several solvers and two user interfaces for defining problem components, all of which reflect a well-defined object model. Hence, the starting point in discussing an object model for *hypre* was the existing one. The goal was to extend this model, making it more precise and consistent, while attempting to reflect the ongoing work of the Equation Solvers Interface (ESI) Forum.

Previously, the *hypre* model was defined only by informal description, convention, and example. The PI worked to formalize and extend the *hypre* object model using the tools and notation defined by the Babel language interoperability project. He also contributed to the initial implementation of several of the classes in this model. A significant advantage of using Babel in the context of *hypre* is that Babel's Scientific Interface Definition Language (SIDL) provides an unambiguous, compilable language for defining an object model, without dictating the implementation language. It proved extremely helpful to have a precise language in which to describe the object model, letting Babel provide (and enforce) the mechanism for implementing the object model in a particular language. As in important side benefit, Babel provides convenient language interoperability, e.g., application codes written in Fortran, C, or C++ can use *hypre* easily, and solver modules written in any of those languages can now be incorporated into the library in a much more straightforward fashion.

While the object *Babel/hypre* model is relatively stable, there is still implementation work to be done, and there are other important issues to resolve. Two of these issues are merging "Babelized" *hypre* with the existing package and responding to the ongoing evolution of the ESI. The *Babel/hypre* work is also being reported in a paper to be presented at the Tenth SIAM Conference on Parallel Processing for Scientific Computing. In this connection, the PI is also organizing a minisymposium at that conference on "Component Architectures for High Performance Scientific Computing."

Summary (continued)

Patrice Turchi is a senior research scientist in the Chemistry and Materials Science (CMS) Directorate at LLNL. He and his co-workers are developing a new, tight-binding-based electronic structure scheme for the prediction of the thermodynamic and physical properties of complex multi-component alloys. The scheme combines molecular dynamics (MD) and monte carlo (MC) simulations so that issues related to different time-scales can be addressed. The energetics which describes the diffusion part of the scheme (MC) is obtained within a real-space electronic structure (ES) method by solving first for the inhomogeneous chemically random alloy in the framework of the coherent potential approximation with a recursion technique, and second for the ordering part of the internal energy with the embedded cluster method and an orbital-peeling technique. Once a new alloy configuration is obtained from MC simulations, the alloy topology is relaxed with MD in the next cycle. The MD-ES-MC loop is iterated until self-consistency is reached and chemical short-range order, energetics, and relaxed topology can be predicted.

Stand-alone sequential codes exist for each of three main components of the MD-ES-MC loops. The MD and ES codes are computationally very demanding, requiring many hours of CPU time for a problem of modest size (about one thousand atoms). During his visit to LLNL, the PI became familiar with these codes, concentrating his efforts on the most expensive step (ES), as implemented in the (sequential) *tbepi code*. This code was first parallelized using OpenMP for the Compaq cluster shared-memory machines at LLNL. Then a distributed-memory MPI version of the code was developed, targeting the IBM ASCI Blue Pacific machine. To take advantage of this architecture a hybrid implementation was developed, with OpenMP-based threads at inner loops and MPI-based problem decomposition at outer loops. A special static load-balancing scheme was developed in order to achieve good performance in the distributed-memory version of *tbepi*. Efficient and scalable performance was demonstrated on up to 2000 atoms on 250 nodes (1000 processors) of ASCI Blue. In addition to the work on *tbepi*, the PI implemented the first MPI-based parallel version of Turchi's MD code *tbmd*.

Development of Object-Oriented Tools for the Numerical Solution of Reactive Flow

Donald W. Schwendeman
Rensselaer Polytechnic Institute

Summary:

Our work developed object-oriented application software, within the Overture framework of codes, for the numerical simulation of high-speed reactive flow. The mathematical model on which the software is based is the reactive Euler equations. The implementation of this model is fairly general and allows for multiple reacting species and reaction rates and a general equation of state with the aim of being able to simulate experimentally observed phenomena in gas or solid explosives. The software is part of the OverBlown package of fluids codes, developed and maintained by Bill Henshaw and the Overture team at CASC. It uses overlapping grids in order to handle general domains and an array class library, developed by Dan Quinlan and the Overture team, which allows parallel processing. The software has been carefully tested for accuracy using existing codes written previously by the author, and it has been used to compute the evolution to detonation of reactive samples subject to various initial conditions and within various confinement geometries. Work on a number of smaller projects has also taken place. These include methods for conservative interpolation on general curvilinear grids, methods for characteristic interpolation, and developing software tools for numerical quadrature on overlapping grids.

A numerical method of solution was implemented for the reactive Euler equations. Within the Overture framework, the numerical method discretizes the equations on a set of overlapping, structured grids that cover the domain of interest. The software package Ogen is used to generate the overlapping grid and provide geometric mapping information for each curvilinear component grid and information concerning the communication of the solution between grids in the overlap region. This information was considered to be given and thus the main task was to implement the numerical method for an individual component grid and then let the existing software, OverBlown, handle the surrounding numerical details (such as interpolation between component grids, the application of boundary conditions, time stepping, graphical interface, and more).

For an individual component grid, an unsplit Godunov-type, finite volume method was used to solve the equations numerically. The approximate Riemann solver is a Roe-type method extended to handle a general equation of state. These ODEs are solved numerically for each grid cell using a Runge-Kutta, variable time step, error control algorithm. This is done to ensure that the (possibly) stiff rate equations are solved sufficiently and accurately.

Summary (continued:

Several example calculations were run to test the software and to study the evolution to detonation of a reactive sample at critical conditions subject to gradients in temperature. The physical motivation for these calculations was to examine the effects of small nonuniformities in the initial state and of the confinement geometry on the severity of the reaction in terms of peak temperatures and pressures. If the gradient parameter for the temperature is too small, then an explosion takes place as a "phase wave" in which a reaction front is born from the initial hot-spot location but travels too fast for significant acoustic information to propagate and strengthen the wave to form a detonation. On the other side, if the gradient parameter is too large, then a local explosion from the initial hot spot does not create a large enough response to initiate significant reaction in the cold neighboring material. In this case, detonation does not occur. In between these extremes, there is a range of values for which the initial gradient leads to a constructive interplay between chemical activity and acoustic propagation resulting in the evolution to a shock-supported detonation. This event leads to high peak temperatures and pressures at the leading edge of the detonation, which are enhanced still further by interactions with a solid confinement

The core elements of the software are now in place and have been tested for accuracy and shown to be useful for numerical simulations. In particular, the software has been used to study the evolution to detonation for a nonuniform reactive sample at critical conditions. While the project has been successful, there are tasks that remain for future work. For example, on the software side, it would be very useful to have an AMR implementation. Work on this aspect of the software is in progress. At present, a Godunov-type numerical flux is implemented with a Roe-type approximate Riemann solver. There are many other choices of numerical flux functions and Riemann solvers that one might use and build into the software as a user option. Additionally, it would be useful to include options for various multiple kinetic reactions, such as chain branching reactions. A long-term goal would be to include diffusive terms in the model equations in order to handle problems involving low-speed flames.

Spatio-Temporal Data Mining of Scientific Trajectory Data

Padhraic Smyth and Scott Gaffney

University of California, Irvine

Summary:

With the increasing availability of massive observational and experimental data sets across a wide variety of scientific disciplines, there is an increasing need to provide scientists with efficient computational tools to explore such data in a systematic manner. For example, techniques such as classification and clustering are now being widely used in astronomy to categorize and organize stellar objects into groups and catalogs, which in turn provide the impetus for scientific hypothesis formation and discovery.

Data-driven exploration of massive spatio-temporal data sets is an area where there is particular need of data mining techniques. Scientists are overwhelmed by the vast quantities of data that simulations, experiments, and observational instruments can produce. Analysis of spatio-temporal data is inherently challenging, yet most current research in data mining is focused on algorithms based on more traditional feature-vector data representations.

Scientists are often not particularly interested in raw grid-level data, but rather in the phenomena and processes which are "driving" the data. In particular, they are often interested in the temporal and spatial evolution of specific "spatially local" structures of interest, e.g., birth-death processes for vortices and interfaces in fluid-flow simulations and experiments, trajectories of extra-tropical cyclones from sea-level pressure data over the Atlantic and Pacific oceans, and sunspot shape and size evolution over time from daily chromospheric images of the Sun. The ability to automatically detect, cluster, and catalog such objects in principle provides an important "data reduction front-end" which can convert 4-d data sets (3 spatial and 1 temporal dimension) on a massive grid to a much more abstract representation of local structures and their evolution. In turn, these higher-level representations provide a general framework and basis for further scientific hypothesis generation and investigation, e.g., investigating correlations between local phenomena (such as storm paths) and global trends (such as temperature changes).

In this work we focused on detecting and clustering *trajectories* of individual objects in massive spatio-temporal data sets. There are two primary technical problems involved. First, the local structures of interest must be detected, characterized, and extracted from the mass of overall data. Second, the evolution (in space and/or time) of these structures needs to be modeled and characterized in a systematic manner if the overall goal of producing a reduced and interpretable description of the data is to be met.

We proposed a general framework for clustering trajectories using probabilistic models of dynamic systems, which allows one to overcome limitations of feature vector-based methods. As such we began looking at various types of dynamic models, for example, autoregressive (AR), moving average (MA), and the more general ARMA model. One of our tasks was to figure out how to simulate direction-focused trajectories from these models in terms of parameter settings. We also looked at learning the parameters of these models given some set of generated data from a known model.

Summary (continued)

An application of trajectory clustering can be found in the clustering of cyclone tracks from meteorological data. We are currently working with a simulated data set for the 1979/1980 winter that gives mean sea-level pressure measurements. We focus on using a bicubic interpolation inside of an iterative scheme to find our minima using a simple gradient descent. First we scan all of the time-sliced images and find all the local minima using a simple sliding neighborhood method. That is, we declare a "pixel" to be at a local minimum if its value is less than all eight of its neighbors. Then we use a simple gradient descent with bicubic interpolation to descend to the point "inside" of the pixel that is at an approximate minimum. This point then gives us our approximate offgrid center of a candidate cyclone. Using the above technique, we processed the data to force all of the grid-based minima to lie in continuous space. We then fed this data into our new tracking software.

We have finished the process of porting our previously developed MATLAB software to our current C++ PC-based platform and, in addition, we have completed the necessary modifications to allow tracking to be carried out using offgrid coordinates. At this time, much of the basic software development has been finished. Much of the future development will be focused on dynamic modelling implementation. We will investigate Kalman filter models.

From here we would like to investigate ways in which the tracking can be integrated into the clustering framework. That is, if we know which cluster an individual belongs to with some probability, then we should be able to more accurately track its future movements. In other words, we believe that instead of two different problems tracking and clustering what we have here is one compound problem that can be solved in an integrated manner.

Algebraic Coarsening Methods for Linear and Nonlinear PDE and Systems

Irad Yavneh
University of California, Los Angeles

Summary:

Algebraic Multigrid (AMG) methods for PDEs were introduced and established in the 1980s. In 1999 Brandt presented a very general approach to algebraic coarsening (including the choice of coarse-grid variables, coarse-grid equations, and inter-grid transfers), which led to an optimal multiscale algorithm for Monte Carlo simulations. The key observation behind this approach, expressed in terms relevant to deterministic PDEs, is the following: Suppose that an appropriate discretization for the fine-grid equations, and some appropriate relaxation method (smoother), have been chosen. Suppose also that an appropriate definition of the coarse-grid variables in terms of fine-grid variables has been selected (e.g., the coarse-grid variables may be a subset of the fine-grid variables, as in classical AMG, or a local average). Define now a compatible relaxation, as an implementation of the above-mentioned relaxation method for the fine-grid variables, modified so as to leave the coarse-grid variables invariant. (For example, in classical AMG this would simply require skipping the variables that coincide with coarse-grid variables.) Then *a general measure of the quality of the set of coarse-grid variable is the convergence rate of the compatible relaxation*. Moreover, fast convergence of the compatible relaxation implies that the dependence of each coarse-grid variable on any other coarse-grid variable decays exponentially (or faster) with the distance between the two, where the distance is measured by coarse-grid mesh-intervals.

The former observation provides us with a tool for gauging the adequacy of our chosen set of coarse-grid variables (and suggests a method for obtaining a set, which is tailored to our needs). The latter implies that a highly accurate coarse-grid approximation to the fine-grid equations can be obtained, where only dependence on a fairly close neighborhood must be explicitly taken into account. It also suggests how such a coarse-grid discretization may he derived by means of solving a local optimization problem. We developed a new compatible relaxation prolongation that seems comparable to classical AMG in its cost of implementation. The new approach offers the potential of generalization to problems where the coarse-grid variables are not (or cannot be) defined as subsets of the fine-grid variables.

As another part of this research, we developed a directly derived coarse grid. To determine the coarse-grid equation for coarse-grid variable, first choose some local fine-grid subdomain, F. We then take a linear combination of the fine-grid equations corresponding to variables belonging to F. This linear combination is itself obviously an exact fine-grid equation. The coarse-grid equation is obtained by ignoring all coefficients corresponding to F-variables. The idea is to choose the coefficients of the linear combination such that the neglected coefficients will be minimal (in some norm), and that the coefficient of the coarse-grid variable will be relatively large so that there will remain an unambiguous correspondence between equations and variables on the coarse grid. The coefficients of the linear combination define a restriction operator. The neglected coefficients of the resulting stencil represent the discrete relative truncation error. We minimize the weighted L_2 norm of the latter, subject to certain constraints. For example, we normally impose that the sum of the neglected coefficients equals zero, which ensures that the coarse-grid approximation of constant functions is exact (since the relative truncation error vanishes in this case).

Summary (continued)

Alternative constraints are also possible, for example. Imposing that all the neglected coefficients be of the same sign, resulting in a coarse-grid operator which satisfies some fine-grid maximum principle. This can be used to obtain monotonic convergence. We implemented this and found the convergence rate was satisfactory, though this approach is expensive. It remains to be seen whether it is advantageous in more general problems.

A central decision in AMG for systems of PDEs is how to choose which equations should participate in the linear combination that makes up each equation in the transformed system. The straightforward approach of including in the linear combination for a coarsened equation all of the fine-grid equations in which it is a dominant variable is unsatisfactory (except for simple problems such as the Cauchy-Reimann equations). We therefore examined a "grey box" version, where we attach a "color" to each type of variable on the fine grid. Then, in the definition of the coarse grid, we apply the usual AMG coarsening process to each color separately (ignoring dependencies on variables of different colors for this process). This approach seems to hold some promise. Indeed, for the Stokes equations with staggered discretization and periodic boundary conditions, this approach can produce an exact decoupling of the equations and the resulting system is solved essentially as fast as the Poisson problem. However, much work is yet needed before a reasonably general "grey box" solver for elliptic systems is achieved. We are currently testing and refining this approach for the Stokes system with boundaries and for linearized Navier-Stokes.

In summary, some new algebraic coarsening methods have been studied, based on ideas related to the condition of fast convergence of compatible relaxation. A generalized AMG approach, where the prolongation operator is derived using compatible relaxation, has been formulated and tested with encouraging results. New methods of direct derivation (i.e., non-Galerkin) have also been introduced and preliminary tests have been made, including for systems of PDEs. Such methods may not be computationally competitive for relatively simple problems, but they can be made quite general.

ISCR

Institute for Scientific Computing Research

Laboratory Directed Research and Development Project Research Summaries

CASC Center for Applied Scientific Computing

SAVAnTS: Scalable Algorithms for Visualization and Analysis of Terascale Science

Mark Duchaineau
Center for Applied Scientific Computing

Summary:

Multi-physics simulation codes on LLNL supercomputers that are vital components of Stockpile Stewardship and other programs may produce dozens of terabytes of data during a given run. Great strides are being made to increase the efficiency and accuracy of the codes by harnessing thousands to tens of thousands of processors using scalable algorithms, but the efficient and accurate post-computation data handling and interactive exploration must also scale efficiently to reach the goal of a productive terascale simulation capability. Our goal is to offer multiresolution data selection and compression algorithms, coupled with optimal data access during interaction, to provide thousand-fold greater interactive performance and hundred-fold more efficient storage than current best practice. Success is measured by the reduction in disk usage, and the increase in interaction rates, with the effectiveness remaining constant or even improving as problem size and machine size grow exponentially.

Our focus is on devising new wavelet transforms and other hierarchies for compression and accelerated access and display of volumetric field data and boundary or contour surfaces. The central principle is output-sensitive calculations: storing and computing exactly enough information given the actual usage patterns of the scientist who is looking at the simulation results. We find that the useful information content of a 3D field such a pressure or density is quite sparse and readily compressed after the application of an appropriate wavelet transform. This results from the property that wavelets tend to automatically find and exploit coherence in both space/time and frequency/scale. Wavelets are well understood for regularly spaced grids filling a complete tensor-product brick of space. The innovations required for large-scale laboratory applications include the extension to highly adaptive or unstructured settings, and to arbitrary surfaces that typically can not be represented as anything resembling a regularly spaced grid. More fundamentally, the work performed should be proportional to the sparse post-transform information content at as many stages as possible of the end-to-end data flow going from simulation to scientist. This leads to a suite of connected optimization problems that we address.

During FY2000 we devised a new type of wavelet for general surfaces that is the first to have bicubic precision while allowing sparse, local transforms using small filters at all steps in the transform. Thus our methods are the first practical high-precision wavelets for use by large-scale simulations for surface storage. Supporting the wavelet transform are new conversion procedures we have devised that automatically shrink-wrap to a semi-regular form the arbitrary-shaped geometry that contains any number of topological handles and connected components. An example is a record-breaking simulation of a Richtmyer-Meshkov instability forming in a shock-tube experiment (for which two members of our LDRD team were co-recipients of several awards, including the Gordon Bell Prize in the Performance Category at the IEEE SuperComputing conference in November, 1999). The complete surface if represented conventionally contains 460 million triangles and occupies 13GB of disk space. The shrink-wrap results depicted allow application of our wavelet transform, which our experiments indicate is likely to reduce the storage to under 260MB, a fifty-fold space reduction. The wavelet transform and coding are at least ten times faster than unstructured surface-mesh compression schemes and previous wavelet surface compression methods based on infinite-width filters.

In FY2000, we plan to increase performance on per-view surface optimization time by two or more orders of magnitude using pre-optimization on subregions, and complete a working prototype for the wavelet surface compression.

Sapphire: Scalable Pattern Recognition for Large-scale Scientific Data Mining

C. Kamath, E. Cantu-Paz, I. Fodor, and N. Tang

Center for Applied Scientific Computing

Summary:

There is a rapidly widening gap between our ability to collect data and our ability to explore, analyze, and understand the data. As a result, useful information is overlooked, and the potential benefits of increased computational and data gathering capabilities are only partially realized. This problem of data overload is becoming a serious impediment to scientific advancement in areas as diverse as counter-proliferation, the Accelerated Strategic Computing Initiative (ASCI), astrophysics, computer security, and climate modeling, where vast amounts of data are collected through observations or simulations. To improve the way in which scientists extract useful information from their data, we are developing a new generation of tools and techniques based on data mining.

Data mining is the semi-automated discovery of patterns, associations, anomalies, and statistically significant structures in data. It consists of two steps: in data pre-processing, we extract high-level features from the data, and in pattern recognition, we use the features to identify and characterize patterns in the data. In this project, we are developing scalable algorithms for the pattern recognition task of classification. Our goal is to improve the performance of these algorithms, without sacrificing accuracy. We are demonstrating these techniques using an astronomy application, namely the detection of radio-emitting galaxies with a bent-double morphology in the FIRST survey.

In FY2000, we focused on three tasks: (a) improving the performance of decision tree algorithms; (b) identifying bent-double galaxies in the FIRST survey; and (c) incorporating our research into software to make it easily accessible to LLNL scientists. In decision trees, we considered oblique trees, where a decision at a node uses a linear combination of features, instead of a single feature. As this is essentially a search in a high dimensional space, we investigated the use of evolutionary algorithms to solve the optimization problem. Our research showed that combining evolutionary algorithms with decision trees resulted in better and faster classifiers. On a data set with 50 features or dimensions, one of our new algorithms, Oblique-ES, was more accurate (79 % vs. 73 %) and four times faster than the current best oblique classifier. Another algorithm, Oblique-GA, which gave the most accurate results (85 %), was twice as fast. In contrast, the traditional axis-parallel tree, while fast, resulted in accuracy (58 %) that was just a bit better than making a random decision.

For the bent-double problem, we focused on galaxies composed of three blobs. Using features from the FIRST catalog, we completed six inner iterations of data mining, improving the features extracted. We completed two outer iterations, increasing the size of the training set from 195 to 495 examples. In the process, we reduced the classification error by 50 percent.

In FY2000, we released the Alpha version of our software, which includes serial, object-oriented versions of decision trees, wavelets and wavelet de-noising techniques, evolutionary algorithms, and support for a parallel infrastructure. In addition, we published 12 papers, 8 conference and workshop presentations, and 3 records of inventions. We co-organized two data mining workshops and actively participated in university collaborations.

In FY2001, we will complete our work on parallel scalable decision tree algorithms and software, complete the detection of bent-double galaxies using the techniques developed, and focus on the use of evolutionary algorithms to improve the performance of neural network algorithms for classification.

ISCR

Institute for Scientific Computing Research

Student Internship Research Summaries

CASC Center for Applied Scientific Computing

Multiresolution Geometry Visibility Determination for ROAM

Lucas Ackerman
Worcester Polytechnic Institute

Summary:

The Realtime Optimally Adapting Mesh algorithm (ROAM) enables Dynamic view-dependent level of detail rendering of multiresolution geometry. In order to reduce rendering overdraw and to allocate triangles only to visible surfaces, it is necessary to have a visibility determination algorithm that operates on multiresolution geometry and scales well to very large data sets.

The algorithm developed for this purpose utilizes an on-demand, dynamic occluder selection process, and can refine occluders as best fits the current viewpoint. The method of scene traversal is based on Frustum Advancement, testing only visible scene portions, and so has a small memory footprint even when dealing with very large geometry sets. To take advantage of frame-to-frame visible set coherence, the algorithm may use eye-cell or occluder-shrinking methods to determine visibility for a region of space around the eye point. The algorithm requires only a coarse octree constructed on base-mesh triangle volumes as a pre-compute step, and so can leverage fast visibility computation of complex scenes without costly computation or storage of potentially-visible-set structures. Finally, the algorithm can also generate a simple quadtree-based Z-buffer, which may be used to check visibility of mobile objects in the scene.

Loop Fusion Optimization for Array Statements with ROSE

Marcel Arndt
University of Bonn, Germany

Summary:

Computing power increased dramatically within the last decade, but performance optimization of numerical applications did not loose any of its importance. Many methods, including expression templates within object-oriented libraries, do a fine job, but are not completely satisfactory, since they are not able to exploit semantic coherences. This is where the C++ preprocessor ROSE, which is being developed within the Overture group at CASC, starts, by applying certain transformations to the source code before compilation. One sort of transformation is based on loop fusion of array statements with respect to cache structures. I developed methods to explore the optimal shape of these optimizations, which are the foundation for a further implementation within ROSE.

The transformed array statements are modeled by C programs, which are generated automatically by several pen- and shellscripts. This allows a hatch-like quasi-exhaustive evaluation of performance subject to more than ten different parameters including degree of dependence between the arrays involved, the amount of loop fusion, the number and sizes of the arrays, and so forth.

For a better understanding and analysis of the results, the data is formatted into clear tables and PostScript files with graphical presentation are generated. For the latter purpose, a C++ program has been developed using Overture.

The results show that the fusion of several loops into one single loop can increase the performance up to 20% because of caching effects, but going too far with loop fusion leads to a code which runs as much as five times slower. We also found that a careful reordering of the statements based on a dependence analysis allows fewer loops and faster codes.

It is planned to implement these transformations with ROSE. My results will serve as a guide for the precise definition of those transformations.

A Least-Squares Finite Element Method for 3-D Neutron Transport Problems

Travis M. Austin
University of Colorado, Boulder
Department of Applied Mathematics

Summary:

The application of a least-squares finite element technique to the 3D neutron transport equation is an interesting alternative to current control volume methods. The method relies on a properly scaled least-squares functional to develop a minimization principle. The scaling is designed to capture the correct behavior of the solution across parameter regimes. The minimization principle provides a solid mathematical framework for building all of the components of a multigrid solution procedure.

Using a finite expansion of spherical harmonics to represent the angular variable leads to a selfadjoint, positive definite system of moment equations. A conjugate gradient method with a block Jacobi preconditioner, which uses multigrid to invert the diagonal elements, is used to approximately solve this system of equations. For nearly all parameter regimes and all diagonal elements, standard multigrid is sufficient as a preconditioner. However, for certain parameter regimes, multigrid performs poorly on the diagonal elements associated with the first-order moments. A new multigrid algorithm is being developed that simultaneously addresses the three first-order moments.

As always, the goal of multigrid is to develop a scheme that has convergence factors independent of discretization size. When the "grad-div" term is dominant, the error reduction factor for the standard multigrid algorithm approaches unity quickly as discretization size approaches zero, which is unacceptable. The problem lies in the "hiding" of some of the oscillatory error terms at the low end of the spectrum. Smoothing works only on eigenvectors at the high end of the spectrum. Usually, these high frequency eigenvectors contain the oscillatory errors. For the given neutron transport operator, other oscillatory errors at the low end of the spectrum are left unattended. When the smoothing does not truly smooth, the coarse-grid correction becomes worthless as a component in the multigrid algorithm.

During five weeks spent at LLNL I worked on development of a 3D code that will tackle this problem of divergence-free error. Britton Chang and Barry Lee of CASC have already shown that the FOSLS approach works very well in two of the three important parameter regimes studied in neutron transport. From our existing 2D experience, we are confident of developing a 3D code that is robust with respect to all regimes encountered in practice.

Djehuty, LLNL's Next Generation Stellar Evolution Code

Brian Ball

Worcester Polytechnic Institute of Technology

Summary:

The Djehuty project is a next generation stellar evolution code. It will be capable of three-dimensional modeling of binary star systems. Since half of all 'stars' exist in binary systems, the Djehuty project will have strategic potential for the stellar simulation community. The Djehuty team at LLNL is comprised of people from many locations, including IGPP, CASC, A and V Divisions, and UC Berkeley. The code is derived from projects at LLNL that have been in existence for over three decades.

One of the main components is ARES, which is a multi-block, regularly connected, 2D and 3D multi-material radiation hydrodynamics code. One of my tasks was to help remove certain functions from the hydrodynamics code which are not applicable for stellar evolution. Another task was the addition and testing of a new equation of state, created by Peter Eggleton (IGPP), in the Djehuty code. Testing was done using a standard shocktube test problem with both the Eggleton equation of state and an ideal gas equation of state at the appropriate environmental conditions. Radiation diffusion was also implemented and tested in the Djehuty code by using a mock idealized spherical NIF capsule. The testing was done in serial and in parallel on the "Compass" Alpha cluster.

Adaptive Methods in Numerical Cosmology

Zackery Belanger
Oakland University

Summary:

The goal of this project is to investigate the use of adaptive mesh refinement (AMR) in the solution of systems of Einstein equations arising in cosmology. Numerical solutions of these systems are increasingly important aids to understanding the nature of singularities characterizing various theoretical models of the origin and evolution of the universe. The solutions of cosmological systems such as the Gowdy T3 model are known to develop localized fine-scale structures as they evolve in time. The efficient solution of such equations requires the ability to employ computational grids with high resolution only where it is needed.

Under the supervision of Xabier Garaizar and Milo Dorr in the Center for Applied Scientific Computing (CASC), I have developed a code to solve the Gowdy T3 system of equations using a Godunov algorithm. The Gowdy system has the form of a variable-speed wave equation with nonlinear source terms. The solutions of this system represent wave amplitudes in a space-time metric. An important issue that I had to address was the treatment of the nonlinear source terms. The development of an accurate and robust uniform grid integrator is the key step in creating an adaptive algorithm using block-structured, locally Cartesian grids. I have begun to incorporate my integrator in an adaptive code using the SAMRAI (Structured Adaptive Mesh Refinement Application Infrastructure) system under development in CASC.

I will continue to develop my adaptive code as part of my Master's thesis at Oakland University. The next goal will be to implement a higher-order version of the current Godunov integrator. I will then use the new code to investigate the adaptive solution of extended models, such as the magnetic and twisted Gowdy systems, which have not yet been numerically simulated, pending the availability of the resolving power of AMR.

Parallel Volume Rendering of Unstructured Data

Janine Bennett
University of California, Davis

Summary:

There is a need for ways to visualize large-scale data sets. Current methods can be very slow and inefficient. By parallelizing portions of HIAC (High Accuracy Volume Renderer), a legacy volume renderer for unstructured data, we can make the visualization of these large-scale data sets much quicker.

The portion of the HIAC code responsible for visibility sorting, polyhedron projection, and the OpenGL calls used to render to the screen is now parallelized using Pthreads and two buffers for inter-thread communication. One thread is used for visibility sorting, one is used for the OpenGL calls, and all remaining threads are used to project the polyhedron. We tested our Pthreads code on a 48-Processor SGI Onyx2 with 8 Infinite Reality pipes running IRIX 6.5.7. The parallelized version of code runs twice as fast as the serial version of the code. Peter Williams, Randy Frank and Deborah Walker were all involved in different stages of the project.

We are investigating ways to further parallelize by processing the input data into several load-balanced screen tiles. For each tile we will apply the threaded code to generate images even faster.

Symmetric Wavelets on Arbitrary Topology

Martin Bertram
University of California, Davis

Summary:

Wavelets are used in lossy and lossless compression schemes for scientific data obtained from flow field simulations on supercomputers. Compression is necessary to store, transmit, and visualize massive volume datasets. Most compression schemes are restricted to data defined on regular, rectilinear grids. For compact representation of isosurfaces and material boundaries of arbitrary topology, however, wavelet compression schemes need to be generalized to handle data defined on polygonal meshes.

We constructed new wavelet basis functions defined on subdivision surfaces that have a polygonal base mesh and generate subgrids of regular topology, such as Catmull-Clark subdivision surfaces. Compared to most wavelet approaches on arbitrary topology, our wavelets generalize tensor product basis functions. As a consequence, we obtain B-spline representations for reconstructed data, except in the neighborhood of a few extraordinary points (with valence not equal to four) in the base mesh.

We implemented two new wavelet transforms that generalize bilinear and bicubic B-splines to mesh domains of arbitrary topology. To obtain compression, we encode the wavelet coefficients that are sparse or have small absolute values using a coding scheme that we developed earlier. The overall compression algorithm is highly efficient and provides surface reconstructions at multiple levels of accuracy.

We will use the new surface compression scheme for efficient extraction and visualization of isosurfaces. We also want to develop similar techniques for compression of volumetric data defined on irregular polyhedral meshes.

Web Voting Application to Assist in Establishing Software Development Standards

Melvina Blackgoat
Northern Arizona University

Summary:

The development of scientific software by and for many parties increasingly requires a mechanism for reaching agreement between the collaborative development groups. There exists no known method of voting on software standards between the Common Component Architecture (CCA) and the Equations Solver Interface (ESI) groups that span several DOE laboratories. Voting and discussion is done during infrequent meetings or through an email based voting system which requires significant amounts of staff time. For collaborative software development to succeed, a more effective voting mechanism is needed to establish software development standards.

We are developing a distributed voting application to facilitate DOE wide convergence on component standards for scientific computing. The web voting application uses a three-tiered architecture: a WWW browser interface implemented using HTML, a middle tier using Java Servlet Pages (JSPs) and additional Java, and a MySQL server as the bottom tier. We began development by creating JSPs that enter forum data into a database. JSPs are a combination of HTML content and behavior coded in Java. The web voting application will allow users to easily access information, make comments and vote on a particular forum in a small amount of time.

This software will be used by the DOE Common Component Architecture (CCA) and Equation Solver Interface (ESI) standards forums to manage voting on software interface standards. Automated web-based voting will simplify multi-lab collaboration and development of new software standards.

Reconstruction of Material Boundaries from Data Sets with Volume Fraction Information

Kathleen S. Bonnell
University of California, Davis

Summary:

There are numerous instances in which it is necessary to reconstruct or track the boundary surfaces (or "interfaces") between multiple materials that arise in simulations. Multi-fluid Eulerian hydrodynamics calculations require geometric approximations of fluid interfaces to form the equations of motion to advance these interfaces correctly over time. This project presents a new algorithm for material boundary interface reconstruction from data sets containing volume fractions.

To solve this problem, we transform the reconstruction problem to a problem that analyzes the dual data set, where each vertex in the dual mesh has an associated barycentric coordinate tuple that represents the fraction of each material present. After constructing a dual tetrahedral mesh from the original mesh, we construct material boundaries by mapping a tetrahedron into barycentric space and calculating the intersections with Voronoi cells in barycentric space. These intersections are mapped back to the original physical space and triangulated to form the boundary surface approximation. This algorithm can be applied to any grid structure and can treat any number of materials per element/vertex.

In typical simulations, the grid cells contain fractional volumetric information for each of the materials. Each cell C of a grid S has an associated tuple ($a1, a2, ..., am$) that represents the portions of each of m materials in the cell, i.e., ai represents the fractional part of material i. We assume that $a1 + a2 + ... + am = 1$. The problem is to find a (crack-free) piecewise two-manifold separating surface approximating the boundary surfaces between the various materials. We consider the dual dataset constructed from the given data set, with each point having the associated tuple. Thus, the boundary surface reconstruction problem reduces to constructing the material interfaces for a grid where each vertex has an associated barycentric coordinate representing the fractional parts of each material at the vertex. We use this "barycentric coordinate field" to approximate the material boundary surfaces.

If we have a data set containing m materials, we process each tetrahedral cell of the grid and map our tetrahedral elements into an m-simplex representing m-dimensional barycentric space. Next, we calculate intersections with the edges of Voronoi cells in the m-simplex. These Voronoi cells represent regions where one material "dominates" the other materials locally. We map these intersections back to the original space and triangulate the resulting points to obtain the boundary.

Concerning future work, we would like to insert a "measure and adjust" feature into the algorithm. Once an initial boundary surface approximation is calculated, we calculate (new) volume fractions for cells directly from this boundary surface. This will enable us to calculate the difference between the original volume fractions and the volume fractions as implied by our initial boundary surface approximation. It is then possible to adjust our material interfaces to minimize the volume fraction deviations. We also plan to extend this algorithm to multidimensional grids.

Line-Implicit Time-Stepping for the Solution of Navier–Stokes Equations on Overset Grids

Lars Carlsson

Chalmers University of Technology, Sweden

Summary:

Boundary layers arise in many fluid dynamics applications. Under popular fully implicit global time-stepping schemes, the solutions are often not resolved in time, and interesting flow phenomena disappear. Instead, a fully explicit method can be used. However, for these methods the time step becomes very small because of the large velocity gradients in the boundary layer. Our goal has been to alleviate the severe time-step restriction without stepping over the interesting temporal behavior. This has been accomplished by using line-implicit time-stepping.

We have implemented a two-dimensional line-implicit time-stepping scheme within the Overture framework. Overset grids have been used to divide the computational domain, and different time-stepping schemes have been used in the advancement of the solution on different component grids. The line-implicit time-stepping is designed to be used in highly stretched near-wall grids, where boundary layers occur.

The Laplacian and other spatial operators in the momentum equations are divided into fast and slow time scales. In the time-stepping, the fast and slow components are treated implicitly and explicitly, respectively. The resulting nonlinear systems of discrete equations are banded. Each system contains the solution along one grid line. The computational cost is of the same order as for a corresponding explicit scheme. The stretching in the direction normal to the implicit direction does not influence the stability of the time integration. The time step is determined only by the eigenvalues associated with the tangential, explicit direction. Hence, the overall time step is greater than in the case of a fully explicit method.

Calculations on applications like two-dimensional wing sections demonstrate how the time-stepping scheme performs in practice.

Algebraic Multigrid based on Element Interpolation

Tim Chartier
University of Colorado, Boulder

Summary:

My research at LLNL is to apply Algebraic Multigrid (AMG) to problems of material elasticity with complex geometry. This work applies to energy research and nuclear weapons stockpile stewardship. To protect our environment, we must understand how these materials behave in storage over long periods. My research plays a role in efforts to create scalable solution algorithms for the PDEs that arise in such applications. In particular, I am researching coloring algorithms for Algebraic Multigrid solvers based on element interpolation (AMGe).

This summer I continued research on selecting a hierarchy of coarse grids that lead to effective intergrid transfer operators by exploiting knowledge of finite element connectivity. AMGe is a practical compromise between purely algebraic multigrid, which makes use only of the final assembled matrix, and methods that require a full geometric description of a PDE problem, such as agglomeration. Element information is frequently available in applications. I continued discussion with LLNL researchers whom are also investigating and testing ideas in this area. I worked closely with Jim Jones on both 2-level and multilevel coloring ideas in AMGe. Our current research focuses on developing a multilevel algorithm that maintains low element complexity, grid complexity, and convergence factors. It is difficult to balance these competing factors. My work has also involved discussions with Panayot Vessilevski, Van Henson, and Rob Falgout. Current research shows promise. Given current advancements, we hope to provide a fully algebraic multilevel AMGe algorithm in 2001.

Parallelization of the HIAC Volume Renderer

Richard Cook
University of California, Davis

Summary:

The HIAC volume renderer is a piece of software that performs volume rendering using unstructured grids. The cells are first sorted using variations of the MPVQ sort by Peter Williams. They are then projected them onto the viewplane using an algorithm by Nelson Max and displayed using OpenGL hardware rendering.

The projection step is the bottleneck in the algorithm. Therefore, in this project, the projection and rendering steps are being parallelized to improve performance on large datasets. It is expected that this project will become the basis for a Master's thesis.

The parallelization of the projection step is done using P-threads under a shared memory processor architecture. Each sorting thread is added to a queue of consumers and one or more threads at the head of the queue are awakened each time the producer finishes a quantum of work. As each finishes, a final rendering thread converts the information stored by the projection threads directly into optimized OpenGL hardware calls.

Under a distributed memory architecture such as that of ASCI White and other supercomputing clusters, distribution of the data set in memory becomes very important. Therefore, the next step will be to load balance each dataset for optimal distributed behavior.

Curvature-based Adaptive Mesh Refinement for Level Set Tracking

Paul Covello
University of California, Davis

Summary:

Level set tracking is used to compute burn tables to be used for detonation shock dynamics codes. The level set represents the detonation shock front, which is assumed to propagate like an optical wavefront and thereby follows an eikonal equation. The advantage of using burn tables is that they greatly speed up and simplify the detonation shock dynamics code.

Algorithms to track level sets exist already. Given a grid, an initial level set, and a speed of detonation shock propagation, these numerical codes predict how the level set will propagate through the rest of the given domain. The advantages of contemporary methods are speed and versatility with unstructured grids. The disadvantage is that when the curvature of the level set is high, the accuracy of these algorithms falls off. There are two ways to remedy this: one is to have higher resolution, or more discrete points, where the curvature of the level set is high; the second is to provide more information to the grid node being updated about the local grid geometry.

I have developed two different algorithms pursuing these strategies. The first performs adaptive mesh refinement in regions of high curvature, and the second increases the discrete scope. Results show an increase in accuracy, sometimes attaining second order. The expected price was computational speed, sometimes taking up to two orders of magnitude longer in execution time. On balance, these algorithms would be appropriate for complex multiscale problems where obtaining similar high accuracy with conventional nonadaptive methods would require several orders of magnitude more execution time.

The results were from two-dimensional problems, both orthogonal and non-orthogonal, and structured, orthogonal three-dimensional problems. The next goal is to demonstrate both approaches on three-dimensional simulations involving both structured and unstructured grids.

Computational Modeling of a Solid Fuel Rocket Booster Using Overture

Nathan Crane

University of Illinois, Urbana-Champaign

Summary:

A Solid Rocket Booster (SRB) is an extremely complex system. Experimental data obtained to measure the performance of a prototype motor or determine the reason for failure of an SRB is both difficult and expensive to obtain. A good computational simulation of an SRB could help alert designers to potential problems without constructing the actual rocket, thus reducing the number of design iterations required. As the solid rocket fuel burns inside of the SRB, complex moving gas/solid interfaces are produced. This geometry is difficult to model with standard grid generation techniques. Overture is an overlapping grid software environment developed at LLNL and designed to combine the numerical and computational advantages of bulk structure with the geometric versatility of unstructured grids. The objective of my research is to investigate the ability of Overture overlapping grids to represent the complex moving geometry inside a burning rocket.

First, a 2D axisymmetric model of a SRB was produced. Incompressible steady state and transient compressible fluid flow equations were solved on the static 2D mesh. The effect of a moving geometry in 2D was also investigated.

Next a 3D computational model representing the full 3D geometry of the inside of a space shuttle SRB was produced. The model represented the flow space inside of a SRB with inflow of hot gas on the void boundary. As the rocket fuel burns the void space in the rocket will expand. The fuel-void interface velocity at a given point is a function of several variables. Propellant burn rate is a function of local pressure, flow velocity, and exposed surface area. In addition, structural deformation of the solid fuel can alter the shape of the void. The size of the 3D model produced (approximately one million grid points) makes time required for solution of the flow equations prohibitive. A grid regression scheme based only on exposed propellant burn area was implemented.

Overture-style overlapping grids proved to be very useful for modeling the complex geometry that may occur during the operation of a SRB.

Large Eddy Simulation of Rayleigh–Taylor Instability using the Arbitrary Lagrangian–Eulerian Method

Rebecca Mattson Darlington
University of California, Davis

Summary:

This research addresses the application of large eddy simulations (LES) to Arbitrary Lagrangian Eulerian (ALE) simulations of Rayleigh–Taylor instability. Large eddy simulation is a popular methodology for dispensing with the artifacts of turbulence modeling at the larger dynamical scales of the flow.

First, ALE simulations of a simplified Rayleigh–Taylor instability are produced. The advantages of ALE over purely Eulerian simulations are shown. Next, the behavior of the LES is examined in a higher-fidelity ALE simulation of the Rayleigh–Taylor instability. The effects of eddy viscosity and stochastic backscatter are examined. The LES is also coupled with ALE to increase grid resolution in areas where it is needed.

Finally, the methods studied above are applied to two sets of experimental simulations. In these simulations, ALE allows the mesh to follow expanding experimental targets, while LES represents the effect of unresolved instability modes.

Domain Decomposition for a Periodic Scattering Problem

Mike Flanagan
Texas A&M University

Summary:

I am pursuing domain-decomposed iterative techniques for wave-Helmholtz problems in two-dimensional periodic layered media. Solution data is passed between adjacent domains via a Dirichlet-Neumann operator. The goals are to establish convergence and determine means to improve convergence rate.

The domain is rectangular with periodicity in the x-direction. The layered structure exhibits parallelism since each structure is identical except for the incoming and outgoing wave data. The discretization on each subdomain is by finite elements with bilinear nodal basis elements. The algorithm employs a direct solver on each subdomain. In each iteration information passes to the next domain (above and below) by updating the outgoing wave and the incoming wave for each respective domain. This process was successful for four layers and this was the main test for working on shared memory systems (4 nodes), but as attempts were made to increase the number of layers, convergence was unreliable. Error decreased rapidly, but then started to diverge. Resonances are suspected.

A rigorous analysis of the mathematics is currently being undertaken. Simultaneously, we are exploring FFT-based discretization techniques to explicitly take advantage of the periodicity.

Data Mining with Sparse Grids

Jochen Garcke
University of Bonn, Germany

Summary:

Data mining is the process of finding hidden patterns, relations, and trends in large data sets. It plays an increasing role in commerce and science. Recently a new data mining technique for the classification problem was introduced by Garcke et al. The current work focused on expanding its capability and applicability for real data sets.

The technique is based on the regularization network approach but, in contrast to the other methods which employ ansatz functions associated with data points, we use a grid in a high-dimensional feature space for the minimization process.

To cope with the curse of dimensionality, we employ sparse grids. Thus, only $O(h \wedge(-1)n \wedge(d-1)))$ instead of $O(h \wedge(-d))$ grid points and unknowns are involved. Here d denotes the dimension of the feature space and $h = 2 \wedge(-n)$ gives the mesh size. To be precise, we use the sparse grid combination technique where the classification problem is discretized and solved on a certain sequence of conventional grids with uniform mesh sizes in each coordinate direction. The sparse grid solution is then obtained from the solutions on these different grids by linear combination. In contrast to other sparse grid techniques, the combination method is simpler to use and can be parallelized in a natural and straightforward way.

We looked into the possibility of applying this approach in the context of dimensional reduction, i.e., eliminating nonessential attributes which did not enhance classification. The results achieved on some test data sets look promising so far. Further we worked on changing the underlying discretization process, moving away from tensor-product based basis functions in d-dimensions to basis functions having d-dimensional simplexes as supports. This will allow the treatment of data sets with a few more attributes than with the old implementation.

Accelerating a Global Ocean Circulation Model

Aaron Herrnstein
University of California, Davis

Summary:

The climate and carbon cycle modeling group at LLNL uses models of ocean circulation to help in the study of human-induced changes on climate and of methods for potentially reducing those changes. As part of this research, we have begun working with a new model developed at LANL, the Parallel Ocean Program, or POP.

Existing ocean models, including the model developed at LLNL, use various numerical techniques to accelerate the convergence of the model solution to an equilibrium state. Two of these techniques are timestep-splitting and deep-ocean acceleration. Both of these techniques involve the rescaling/redefinition of time within particular model equations or geographic regions.

Without the use of these techniques the models take unreasonably long amounts of time to reach an equilibrium state, often more than two months of time on a massively parallel computer. Aggressive use of acceleration can reduce the time to equilibrium by more than a factor of 50. However, these techniques are largely untested in the POP model, so the objective of this work is to perform such testing.

To test the applicability of these acceleration techniques in POP, we have made a series of model simulations with different combinations of acceleration. The results of these runs will be compared to determine the specific impacts of these combinations on the equilibrium solutions.

The runs are accelerated toward equilibrium solutions by factors from 5 to nearly 100 faster than an unaccelerated model. More than 30,000 processor hours of an IBM SP computer were used—this corresponds to running a high-power desktop computer for more than 3 years and demonstrates the need for such acceleration techniques. Even so, more run time is still needed to bring one of the simulations to equilibrium.

Analysis of the data generated by these runs will be ongoing, concentrating on the potential distortions to the seasonal cycle and the impacts on the large-scale circulation, such as the poleward heat-transfer and strength of major currents. Also, more simulations will likely be needed to determine the best operational mode for the model.

Developing a Higher Order Godunov Method for Lagrangian Based Shock Hydrodynamics

Charles Hindman
University of Colorado

Summary:

Computational shock hydrodynamics is a difficult problem occuring in many fields, including aerodynamics, astrophysics, material deformations, and weapons design. We are investigating a new methodology for solving shock hydrodynamics problems by coupling a higher order Godunov method with a cell-centered Lagrangian scheme. The combination of the two allows the creation of a code that exactly conserves energy and momentum and does not require the use of an explicit artificial viscosity or hourglass filter, unlike typical staggered grid approaches.

Using the framework of A++ and Overture, software toolkits developed at Lawrence Livermore National Laboratory, a non-staggered grid, Lagrangian-based higher order Godunov solver was developed for the solution of the Euler equations in one and two dimensions. Various algorithmic options were included in this code, including velocity and energy updating methods and viscous damping, for testing and comparison purposes. The ability to set up different initial conditions and boundary conditions for various problems was included.

This code was tested on various common verification problems, including the Sod shock tube problem, the Woodward–Colella interacting blast wave problem, the 2-D Sedov problem, and the Noh problem in one and two dimensions. The results matched those reported in the literature for all cases except the 2-D Noh problem, which experienced significant mesh entanglement along the *x*- and *y*-axes. However, this is known to be typical of staggered grid Lagrangian methods as well. The next step will include comparison testing and speed-up analysis.

Adding Embedded Boundary Capabilities to SAMRAI

Jason Hunt
University of Michigan

Summary:

Engineering applications are often multi-scale. Instead of resolving the entire domain to the smallest scale, structured adaptive mesh refinement (SAMR) refines the mesh local to areas requiring higher resolution. LLNL's SAMRAI framework uses object-oriented techniques to provide a parallel SAMR framework for application developers. Engineering applications also involve complex geometry within the domain of interest and many interesting flow phenomena are driven by interactions with complex geometry. Therefore, it is very desirable to combine SAMR with complex geometry representation.

The current version of SAMRAI and a previous implementation of the embedded boundary technique, redistribution, were used as a starting point. The previous embedded boundary implementation was altered to be compatible with the current version of SAMRAI, and new data structures were introduced to enable more efficient algorithms to work with the embedded boundaries. We achieved embedded boundary capabilities for rectangular coordinates in two and three space dimensions. The code currently supports simple analytic geometry. With the framework that SAMRAI provides, we then combined the embedded boundary technique with AMR for rectangular coordinates in two space dimensions. Testing of the code was also performed to address conservation violations.

The re-redistribution technique, which ensures conservation near embedded boundaries that pass through differing levels of refinement, needs to be implemented, and work to enable parallel computations is required. In addition to extending the current work in rectangular coordinates to three space dimensions, there is interest to implement cylindrical coordinates in two space dimensions. Finally, the goal is to have a SAMR code that can handle general complex geometries.

The Validity of Paraxial Approximations in the Simulation of Laser-Plasma Interactions

E. McKay Hyde
California Institute of Technology

Summary:

High-intensity lasers such as those used in inertial confinement fusion produce high-density plasmas, which interact with the propagating light. Solving the Helmholtz equation to compute the laser scattering induced by variations in the plasma density remains too difficult for large computational domains. Hence, various paraxial approximations are often employed because of their efficiency. In this work we sought to establish the domain of validity of these paraxial approximations.

We compared the light intensities computed by solving the Helmholtz equation and five different paraxial equations on small computational domains and for a variety of density profiles. The Helmholtz equation was solved using a variant on the method recently described by Bruno and Sei and the paraxial equations were solved using the implementation described by Dorr and Garaizer of CASC. The density profiles were of two types: (a) small, analytically described combinations of peaks and valleys in a constant background density; and (b) small sections of density profiles computed in actual laser-plasma simulations.

Although it is difficult or impossible to describe a complete set of criteria by which one can determine if a given paraxial equation will be valid, we did observe a general trend. For background densities at and below 0.25 nc (nc is the critical density), the Helmholtz and paraxial equations all agree rather well. However, for background densities at and above 0.4 nc, we observe significant disagreement. Furthermore, no single paraxial approximation emerges as the best choice in all cases. Some features of the solution are captured well by one paraxial equation which are not captured well by another, and vice versa.

Analytical and computational comparison of the various paraxial equations continues. We hope to find ways to improve or at least predict solution accuracy without sacrificing too much efficiency.

Improved Parallel and Serial ILU Preconditioning

David Hysom
Old Dominion University

Summary:

Many scientific codes ultimately devolve to the solution of large, sparse systems of linear equations, $Ax = b$. The linear systems, whose solutions may consume a vast majority of the execution time, are frequently solved using preconditioned Krylov methods. Popular incomplete factorization (ILU) preconditioning, although popular and comparatively robust, is in general difficult to parallelize. Additionally, in many cases it is unclear how best to select the most effective runtime parameters (e.g., the fill level for ILU(k) or the dropping threshold for ILUT). We are investigating various shared-memory parallelization approaches and means of increasing preconditioner effectiveness while reducing preconditioner size.

A literature search reveals prevalent beliefs that preconditioners should be relatively small (have approximately the same number of nonzeros as the problem matrix, A), and that preconditioner formation (factorization) and application (triangular solves) should take little time, compared to the Krylov method. There is increasing evidence, however, that larger preconditioners (i.e., wherein more fill-in is allowed) are most effective at minimizing run time for certain classes of problems. One question investigated, therefore, was how to maintain the effectiveness that comes with increased (higher level) fill, while simultaneously reducing preconditioner size.

We experimented with several sparsification strategies: dropping small values from A before factorization; dropping small values from the factors; dropping with and without row-scaling, etc. For the systems studied, we found that dropping small values could reduce solution time by 10% or more, and dropping small values from A was somewhat more effective than dropping values from the factor. It remains unclear, however, how best to select dropping thresholds for production-code applications.

A second means of reducing preconditioner size is to compute and apply the preconditioner in single precision arithmetic. This strategy has the effect of reducing preconditioner size in terms of memory requirements, as opposed to nonzeros. Although slightly increasing iteration counts, single precision preconditioners were found to reduce run time by up to 20%.

I previously developed and published results for a parallel ILU (PILU) algorithm implemented using MPI, and showed scalability to several hundred processors. While at LLNL I became interested in a medium-scale, shared memory implementation. The motivation was to reduce solution time for a specific set of problems (e.g., 2D CFD filament problems of Petri Fast) produced by the LLNL's Overture code. These problems contain up to 400,000 unknowns, and were being solved sequentially, since Overture's current implementation is sequential. Using OpenMP, I implemented threaded versions of Block Jacobi and PILU preconditioners, and CG and BICGSTAB Krylov methods. Comparisons on the Compass Alpha cluster are indicate high promise for PILU. The resulting code is expected to be incorporated into a future release of Overture through the Hypre library.

Constrained ALE Grid Dynamics

Ana Iontcheva
University of California, Davis

Summary:

In transient structural dynamics simulations at LLNL, objects are typically modeled with unstructured hexahedral meshes. As the objects undergo compressible motion, the grid distorts according to predefined rules so that the accurate answers can be computed for the field equations while maintaining regularity in the grid connectivity. Often, the direction of portions of the grid must be forcefully constrained so that they do not venture into spatial regimes where they violate certain physical principles. For example, in contact-impact problems grid points from each object are constrained to remain in their physical space. The imposition of these constraints is performed using a Lagrange multiplier method and the grid points are moved via the solution of a matrix equation.

During the summer, the two-dimensional motion of a continuum over a solid object was studied. The continuum was modeled with a two-dimensional viscous Burgers' equation in Lagrangian coordinates. The grid motion was caclulated using the velocity of the medium. The original grid was defined using the True-Grid package. As predicted, grid distortion occurred very rapidly. Future work includes the use of a viscous grid motion equation with Lagrange multipliers to enforce the grid motion around the solid object.

Parallel Implementation of Mortar Finite Element Method in Three Dimensions with Multigrid Preconditioning

Chisup Kim
Texas A&M University

Summary:

The mortar finite element method is a nonoverlapping, nonconforming domain decomposition method in which each subdomain can be meshed independently. This provides a natural setting for parallel implementation. The resulting algebraic system, be it from parallel or serial implementation, requires efficient iterative methods and for this purpose, we construct a dual basis multiplier space and a multigrid preconditioner.

The continuity of the solution across the generally non-matching subdomain interfaces is imposed in a weak sense by using the multiplier spaces. At each such interface, a mass matrix has to be inverted. Our dual basis multiplier space consists of discontinuous functions and can be constructed in a simple fashion. Moreover, it produces diagonal mass matrices, which is essential to efficient computation of three-dimensional problems.

The mortar problem can be formulated as a nonconforming finite element method, which leads to a symmetric positive definite problem. The multilevel mortar finite element spaces are constructed based on successively refined meshes starting from a coarse one, and are nonnested. The multigrid algorithm considered is an extension of the two-dimensional case analyzed recently by Gopalakrishnan and Pasciak, which in turn is an application of the BPX method. We use the dual basis multipliers to further improve the efficiency. This produces a uniform preconditioner; that is the condition number of the preconditioned system is bounded independent of the problem size.

The parallelization was done using OpenMP, which is an interface for parallel programming using shared memory. This machine provides shared memory in the hardware. The memory is physically distributed across processors, but is accessible by all the processors.

In algebraic multigrid, one starts with a fine grid, matching or nonmatching, and obtains coarse grids by applying successively a coarsening strategy, e.g., element agglomeration. When coarsening independently in parallel in each subdomain, the resulting coarse grid will not in general match across subdomain interfaces. Furthermore, the elements at such interfaces will not necessarily have regular shapes such as triangles or rectangles. Mortar finite element method with dual basis multiplier spaces in this context will be studied and implemented to be run on a parallel machine.

Using Sapphire to Mine Astronomical Data

Imelda Kirby
University of Washington

Summary:

Data collection threatens to overwhelm our capacity to organize and review it for scientific information and inference. Project Sapphire in the Center for Applied Scientific Computing at LLNL is developing new techniques for analyzing massive quantities of data to extract features that are useful for scientists engaged in the research.

My project in data mining for astronomical data began with installing IRAF, and other software for astronomical images. I continued last summer's search for different automated ways of detecting asteroids in terabytes of astronomical data, this time focusing on SDSS (Sloan Digital Sky Survey) data. I studied the earlier successful application of decision trees to the classification of bent-double galaxies of the FIRST data set, while trying to develop some image processing techniques that would improve the results of the decision tree.

Algebraic Multigrid Homogenization

Stephan Knapek
University of Bonn, Germany

Summary:

Solutions for problems that model locally strongly varying phenomena on a micro-scale level require that all length scales present in the problem be resolved. However, due to storage requirements and numerical complexity, the grid for numerical simulation cannot be chosen fine enough to meet this requirement. Fortunately, in many practical applications the fine-scale details of the solution are not of interest, but only a coarse-scale solution is sought. Therefore, it is sufficient to work with averaged equations. The first step towards an accurate numerical treatment of these problems is the determination of an approximate mathematical model that captures the influence of the unresolved fine scales of the medium, and hence effectively disconnects the mesh size of the computational grid from the size of the heterogeneities. This "upscaling" or homogenization procedure results in equations with so-called effective coefficients that vary on a coarse scale. We investigate the use of algebraic multigrid coarse-grid operator constructions for upscaling and homogenization by interpreting Galerkin coarsening methods as discrete homogenization methods.

Our approach for the computation of a coarse-scale operator and an approximation to the effective diffusion coefficient is based on methods that have been used for some time in robust multigrid algorithms. Specifically, the coarse-grid operators that are constructed with matrix-dependent prolongations by means of the Galerkin approximation may be viewed as discrete homogenized operators. From the coarse-grid limit, operator approximations of the effective diffusion coefficient can be determined. However, we showed that a straightforward interpretation of the standard AMG coarse-grid operators does not necessarily lead to a reasonable homogenized operator. This is due to the choice of coarse grid points in the AMG method. Our current research is on constraints to place into the AMG-coarsening to lead to reasonable homogenized operators.

A Parallel Three-dimensional Magnetohydrodynamics Solver

Joseph Koning
University of California, Davis
Department of Applied Science

Summary:

The simulation of magnetohydrodynamics (MHD) requires the solution of fluid dynamics coupled with electrodynamics. An interesting problem in this field is a compact toroid moving through a plasma contained within a tokamak fusion device. An idealization of this problem is a superconducting sphere moving through a plasma. The sphere's motion will excite waves in the plasma. The energy loss through the interaction of the sphere with the plasma and magnetic field will result in the sphere slowing down and eventually stopping. This model system has been treated analytically by Newcomb. The superconducting sphere in an ideal plasma serves as a proof of concept for the problem of a compact toroid interacting with a tokamak plasma. This problem is three-dimensional and includes the fluid dynamical phenomenon of shocks, as well as electrodynamic phenomena such as Alfven and magnetoaucoustic waves.

I use an Eulerian fluid dynamics method coupled with a Vector Finite Element (VFE) method to simulate the sphere-plasma interaction. The Eulerian method treats the fluid description of the plasma and the VFE method treats the electric and magnetic fields. The vector basis functions of the VFE method can be used that maintain normal continuity for the magnetic field and tangential continuity for the electric field while maintaining zero divergence of the magnetic flux density.

Implementing the method involves combining a modem object-oriented fluid dynamics code with the object-oriented VFE code. Currently the parallel framework PETSc is used to integrate the VFE method that has been combined with the fluid dynamics code. A cross product operator has been implemented to accomplish this combination. The current framework has the ability to construct cross and dot products as well as the curl and div differential operators, for the vector basis, and grad operators for the scalar basis functions. Some preliminary fluid dynamics simulations of a sphere propagating through an uncharged fluid have been run. Future work will involve validation of the current framework and simulations of the sphere propagating through the charged fluid.

Development of Flexible Smoothers for Overture

Frank Koster
University of Bonn

Summary:

Smoothing operations are required at various places within LLNL's Overture code, for example in a multigrid scheme or as part of elliptic grid generation. Therefore, a class of flexible smoother functions should be designed for reuse in all the various tasks.

The objects a smoother acts upon are the coefficient matrix, the right-hand side, and the solution. These objects have to provide the smoother with functions for the defect calculation, the (optional) treatment of boundary conditions, and various special elements, e.g. the diagonal of the coefficient matrix. Through such an object-oriented approach, the smoother attains the desired flexibility.

In our design these functions are members of the supplementary class "CoefficientMatrix" which hides all details to the smoother. Presently, the class "CoefficientMatrix" can deal with second and fourth order operators in two and three dimensions, with general, cross, and constant stencils. Boundary treatment is presently provided for second order operators only.

For the smoother functions, there are two interfaces: "realMappedGridFunctions" and the low level "realArray." Furthermore, there are three ways of using the smoothers: "realMappedGridFunction" with inherited topology and boundary treatment; and "realArray" with or without inherited topology, and no boundary treatment. This provides the flexibility required for current and future adaptive mesh refinement extensions of Overture.

We will complete the present design with boundary conditions for fourth order operators and line smoothers.

Superresolution of Buried Objects in Layered Media by Near-Field Electromagnetic Imaging

Sean Lehman
University of California, Davis

Summary:

In the use of non-invasive waves of either electromagnetic or acoustic origin to probe layered media in near-field conditions, few researchers outside of optical microscopists have taken advantage of the evanescent part of the scattered field to enhance resolution. In this project, we have proposed and are developing an imaging technique to be used in near-field environments. Our technique achieves resolution beyond the diffraction limit, or equivalently "superresolution," by including the evanescent part of the field backscattered from objects buried in the medium.

Tomographic imaging is a collection of techniques to reconstruct images of the unknown internal structure of an object from fields transmitted through, and/or reflected from it. There are two widely used tomographic techniques: projection tomography and plane-to-plane backpropagation. The most widely used diffraction tomography technique is resolution-limited to approximately a half wavelength. We study the forward scattering process in order to develop a new diffraction tomography imaging technique which achieves resolution beyond the classical limit. We proved a new theorem which explains why the diffraction tomography method is resolution limited. We then derived a total field scattering relation, which includes both propagating and evanescent field components. We are developing a new reconstruction algorithm based upon the total field scattering relation. This full-field tomographic reconstruction algorithm includes both propagating and evanescent field components. We present reconstruction results from both simulated and real wide-band radar data which demonstrate that our new algorithm surpasses the resolution of most current techniques.

Alternating Tridiagonal Solvers for the Anisotropic Diffusion Equation Spatially Discretized by Linear Triangular Finite Elements

Linh Lieu
University of California, Davis

Summary:

Alternating Tridiagonal Solver methods here refer to generalizations of alternating-direction-implicit (ADI) methods taking advantage of direct tridiagonal system solution without regard to notions of spatial direction. Stable ATS methods for evolution/solution of the anisotropic diffusion equation on spatially nonuniform grids could significantly reduce execution times. Our goal was to test ATS methods on logically uniform, spatially nonuniform grids, using linear triangular finite-element (LTFE) spatial discretization.

ATS, Crank Nicolson (CN), and Backward Euler (BE) methods were implemented in our code as evolvers. Also, ATS was implemented as a solver (evolution to steady state). Whenever ATS was used, the spatial operator was additively split into tridiagonal matrices, each corresponding to a logical direction. For LTFEs on logically uniform triangular meshes, this splits the spatial operator into three tridiagonal matrices. These matrices were used in a Douglas-Gunn scheme to advance the diffusion equation by some small (evolver) or large (solver) time step. The Doss-Miller prescription was used for time step adaptation.

As an evolver, ATS was used with the FE lumped mass matrix. However the use of a consistent mass matrix with CN or BE improved accuracy. When stable, the ATS method is algorithmically efficient. However, stability has been achieved only on spatially uniform grids, and for LTFEs, only when all interior angles of the triangles are greater than 45 degrees. So far the method is conditionally stable on spatially nonuniform grids, so no general utility has been found. However, lumped mass ATS evolution on spatially uniform grids is sufficiently fast and accurate that it is worth using on appropriate large problems.

An Evolutionary Algorithm Library for General-purpose Optimization

David W. Littau
University of Minnesota

Summary:

Evolutionary algorithms are useful for solving optimization problems. They can be applied to any black box function with parameters which, when varied, produce outputs of varying quality. Unlike many traditional optimization techniques, no other information about the function is required.

Many techniques exist in the evolutionary algorithm hierarchy. It is desirable to have a general-purpose library of specialized functions that can be combined to both reproduce existing evolutionary algorithms and to experiment with new ones. Since evolutionary algorithms are inherently parallel, parallel techniques should be included.

The library was implemented in C++. Pure virtual base classes were used for the crossover, mutation, and selection classes to provide a common interface among the various techniques available to perform each operation. Inheritance from these base classes enforced the common interface, allowing the seamless integration of the different methods into a wrapper class. The wrapper class results in a framework in which a scientist who is interested in optimizing a particular function can do so without needing any knowledge of how evolutionary algorithms operate. A comprehensive collection of classes is available.

Experiments were performed with the code library, and some improvements in existing techniques were realized. The research with the library is ongoing. Furthermore, the library will be used to enhance the performance of some tools that are currently being used on other research projects.

Automated Scalability Performance Analysis of MPI Programs

Michael McCracken
Pennsylvania State University

Summary:

Scalability performance analysis of message passing parallel programs is a difficult task, especially as the size and complexity of programs and their respective trace files increases. Traditional forms of trace data analysis, such as visualization, can become difficult to use when analyzing very large parallel jobs, such as those being run by researchers at Livermore using the ASCI platforms. Scalability is an important concern for these large parallel codes; however, measuring it is difficult because one must analyze data from numerous performance experiments. Tools that give summary-style and specific prescriptive information are now necessary to understand the scaling behavior of large parallel programs.

We have developed MPSET, an innovative performance tool for MPI programs. This tool gives the programmer specific prescriptive information on the scaling behavior of components of their application, enabling analysis of large systems with significantly reduced effort. MPSET examines multiple trace files from runs at various numbers of tasks, and then shows MPI communication operations that scale poorly. MPSET calculates the mean duration of each MPI library call at each scale, and can then fit a curve for each call across the scales, highlighting curves that denote poor scalability or irregular performance. This allows the programmer to better focus their optimization effort. Results from the ASCI Compact benchmarks and the NAS parallel benchmarks demonstrate that MPSET can significantly simplify exploring the cause of scalability problems and, thus, improve the scalability of MPI applications.

Further development of MPSET is expected, with the possibility of tying it to a user-interface infrastructure in development at CASC, as well as adding useful analysis features to MPSET, such as automatic highlighting of MPI functions that scale poorly and possible work on automatic load imbalance detection.

Reengineering and Extending the DataFoundry Applet

Jason V. Morgan
University of Utah

Summary:

Prototype software is often developed without fully incorporating modern software engineering techniques. However, before the software is put into production use, the code is usually reengineered to utilize those techniques in an effort to reduce future maintenance costs. DataFoundry's Java applet, the graphical front-end to a collection of genomics data and tools, is starting to be used in genomics research and needs to be made easier to extend and maintain. New features also require introduction to improve the software's usefulness.

Originally, two large Java classes, DF_applet and GeneticsFrame, provided the bulk of the DataFoundry interface functionality. DF_applet defined the components required to form a query and served as the starting point for the interface. GeneticsFrame defined the interface for viewing and manipulating the results from a query.

New queries, with different input and coding requirements, were difficult to add to DF_applet. To address this problem, DF_applet was split into smaller classes. One of these manages the query interfaces, while the others provide the specific functionality required for each query.

Adding new features to GeneticsFrame was difficult and error-prone because of both its size and the complex interaction of internal components. This problem was addressed by identifying similar data structures and methods, then creating several generic classes to provide the same functionality. These new classes promote significant code reuse and thus are much easier to maintain. A second change, one that made extending the applet easier, was the addition of code supporting automatic menu creation through a distinguished class hierarchy.

The most prominent new feature added to the interface was the ability to view the results of a batch blast run. This required defining BatchBlastPanel, the query starting point, and BatchResultsFrame, which displays a high-level view of the results. The frame also provides several new data manipulation capabilities, including the ability to email data in a spreadsheet compatible format.

Reducing future maintenance costs required more than just good engineering techniques. During reengineering, the code was made more readable, and extensive documentation was written. Near the end of the summer, additional time was spent writing high-level documentation to guide future programmers.

While the redesign of the applet is complete, the DataFoundry team will continue to add new features in response to user requests.

Software Development for Sapphire

Dave Nault
University of Cincinnati

Summary:

We undertook several software engineering improvements on the Sapphire framework for automated scientific data mining.

We tested the Sapphire Basic Toolbox Classes, creating a new set of test programs and making changes to enhance usability and documentation.

We optimized the Histogram Classes by removed dependency on the Standard Template Library, creating a new base class to manage parallel processing and interval manipulation, and unifying the histogram class hierarchy.

We also created new classes to handle array indexing, resizing, and resource management.

We updated the Decision Tree code, changing the way nominal data is represented, now using numbers instead of strings. Five new splitting criteria classes were added, along with new code to use the new histogram classes. Together these changes lead to a 100% performance gain.

Finally, a new data viewer was created. This included defining a file format for images, modifying the FITS viewer to read files in the new format, and writing some glue code that allows any program to display a data object by writing it to a file and launching the viewer application.

Overlapping Grid Interpolation using the P++ Array Library

Stefan Nilsson
Chalmers University of Technology, Sweden

Summary:

A key ingredient when solving Partial Differential Equations (PDEs) using overlapping grid techniques is the interpolation of boundary values between component grids. As it is an essentially "unstructured" operation, its implementation on a parallel computer will differ substantially from that of the difference operators usually employed at the inner grid points.

We have implemented a small subset of the possible overlapping grid operators using the parallel array library P++. The serial version of P++ (A++) is presently used in the Overture project for overlapping grid computations. To implement the interpolation algorithm we had to temporarily leave the P++ framework and code the necessary message passing explicitly using MPI function calls.

Subgrid Scale Reaction Modeling for Application to Turbulent Reacting Flows

Diem-Phuong Nguyen
University of Utah

Summary:

My research involves developing a computational fluid dynamics (CFD) code that calculates both turbulent reacting flows and complex chemical kinetics. Practical reacting flow simulations are accomplished through the introduction of a subgrid scale (SGS) reaction model. My summer research at LLNL involved using Overture to incorporate a SGS reaction model to transport equations, bridging microscopic details to the macroscopic domain.

The Overture framework allows for simulation of physical processes in complex geometry. A general Overture code was written in C++ to solve a system of PDE's on a composite grid using multiple time-stepping techniques. Both Euler and Crank Nicholson time stepping were implemented. Several overture functionalities were exploited, including interactive plotting, time-step control based on the CFL number, and twilight zone forcing which uses the method of analytic solutions to generate an exact solution to the PDE.

Convection-diffusion-reaction equations for multiple species were tested. An SGS reaction model involving the Baldyga two-step reaction mechanism was used to provide the reaction source term in the species transport equations. The Baldyga reaction model calculates the chemical state space of the system given the independent mixing and reaction progress variables. The reaction rates are returned to the transport equations and the new chemical species concentrations are calculated. This coupling of the reaction model to the transport model is achieved via a Newton search algorithm.

The reaction computations were calculated on both single and composite grids. Comparison of the results yielded no difference between the two cases for the same geometry.

I plan to run more detailed flow simulations involving the Baldyga reaction model in Overture for validation purposes. I also plan to couple different and more complex reaction models to the code. Other models include equilibrium and a 6-step heptane mechanism, which includes soot chemistry.
Comparison of the different reaction models can then be made by running the different computations in the Overture framework using the same geometry. I will also test the system in a parallel environment.

2D and 3D Models of Rabbit Sinoatrial Cells Using Overture and CVODE

Chris Oehmen
University of Tennessee, Memphis

Summary:

Over 300,000 Americans die each year from sudden heart failure. Although some aspects of cardiac function have been studied for over one hundred years, the mechanisms by which cardiac fibrillation and defibrillation operate are still uncertain. Many investigators have succeeded in modeling various electrophysiological and metabolic aspects of cardiac cellular function, but more detail and computational power are needed to discern the interplay of cell- and tissue-level processes.

The present implementation of a rabbit sino-atrial node cell is based on the model and code of Demir et al., which was modified to contain additional current features not known until after the time of publication. For instance, the rapid and slow delayed rectifier potassium currents and the rapid and slow components of hyperpolarization current activation were both incorporated in the model. In addition, the model was developed to reflect the mean action potential (AP) characteristics observed by Verheijck et al (1998) for spindle cells.

The LLNL code Overture was used to simulate the cell in 2D and 3D environments. In the extracellular and intracellular spaces, simple diffusion was solved for K+ and Na+. Ca2+ was allowed to diffuse in the extracellular space, while it was uptaken and released from another interior space, the Sarcoplasmic Reticulum (SR), in accordance with the Demir model, as well as buffering in the SR and intracellular spaces.

The concentration gradient across the membrane determined the Nernst potential at all points on the membrane, for which the LLNL code CVODE was used to solve a system of nonlinear ODEs that determined the membrane currents. This membrane current activity was then allowed to interact with the diffusing species in the various spaces.

Overture allowed for rapid implementation and development of the models, as well as 2D, 3D, and movie-style visualization of the results. CVODE provided a robust mode of integration for variables on the membrane surfaces.

The models were in good agreement with those in the literature. Both the 2D and 3D cell models exhibited the critical aspects of rabbit SAN function, including endogenous pacing and the relative dominance of particular membrane currents during different phases of the action potential.

Future developments of the models include the incorporation of electric field calculations based on the charge distributions calculated by the models. This will be an integral part of the final models in understanding the behavior of the ions in the represented physiological systems.

If the models are further validated with respect to the processes of interest, they may be extended into a massively parallel implementation in which many cells can be modeled in 3D, and their tissue-level properties can be studied.

Scalable Domain Decomposition Algorithms for Resolving Contact Surfaces in ALE Computations

Tim Pierce
University of California, Davis

Summary:

In the Lagrangian approach to modeling material media, the overall volume of material under consideration is subdivided into a discrete set of zones, together with their connecting faces, edges, and nodes. The connectivity of the resulting mesh determines which nodes belong to each zone, from which can be deduced which zones contain a given node. This information establishes the proximity relations of all zones.

As the material moves in time under the influence of internal stresses, body forces and boundary conditions, the mesh follows the material, so that a given zone always contains the same material. For an explicit time advancement scheme, the behavior at a given node or zone over a single timestep is dictated only by material in its own and neighboring zones, which can be determined from the mesh connectivity. The timestep is chosen, with reference to the speed at which signals can propagate through the material (typically the sound speed), so that this locality of influence is a valid assumption.

If the model includes more than one discrete body, then the interaction of the bodies must also be considered, and mesh connectivity is no longer sufficient to determine proximity. The bodies may come together or separate, or may slide along each other, each exerting a boundary force on the other. Though the connectivity of the mesh representing each individual body is constant, which zones are adjacent across the contact surface is entirely dynamic and unpredictable.

A standard approach to parallelizing a Lagrangian dynamics code involves dividing the mesh into a number of submeshes, or subdomains, of approximately equal size, and assigning the subdomains to separate processors. A good decomposition minimizes communication, as only neighbor information is required. Furthermore, since the connectivity is constant in time the communication pattern can be set up during an initialization phase, and only the variable data itself updated each cycle.

This simple domain decomposition method of parallelization breaks down at contact surfaces, or at least an additional mechanism must be supplied to: (a) determine proximity relations across contact surfaces; (b) distribute the calculation of the contact forces across the parallel machine; and (c) gather the necessary data from across the machine together on the processor where that part of the surface is to be calculated. The task of calculating contact forces can be divided into two steps: contact detection and contact enforcement. First one must determine if and where two surfaces are in contact. Then one can apply suitable balancing forces at those points.

In a serial environment, the detection problem can always be solved by brute force. The position of each node on one side can be compared to the positions of all faces on the other side, the closest face determined, and then penetration checked for. In a parallel environment, however, even the brute force approach is not available, since the closest face on the other side may be on a different processor. As time changes, the closest face and perhaps the processor assigned to that face may also change. In fact, the parallelization of contact surfaces provides formidable challenges, particularly if the task must be done in a scalable fashion as the mesh size and processor count are increased proportionately.

In this research, we developed a truly scalable solution, implemented it in a major dynamics code (ALE3D), and demonstrated scalability well beyond 1,000 processors. A document describing the design and implementation of the algorithm in great detail is currently under development.

Shrink-wrapping Algorithms for Large Data Visualization

Serban D. Porumbescu
University of California, Davis

Summary:

Scientists are doing simulations that produce terabytes of data output. Ideally, scientists need to interact with and manipulate the data in real time. Unfortunately, contemporary computers have I/O systems that are out of balance with the rate at which they can create data. To achieve our goal of visualizing large amounts of data we need to invent new algorithms that take advantage of powerful wavelet compression schemes.

Given an unstructured surface, shrink-wrapping produces the topologically equivalent semi-structured mesh needed for wavelet compression. The algorithm begins with a coarse base mesh and repeatedly subdivides, smoothes, and snaps (moves towards the desired surface) until the mesh parameterization is complete. Our principal improvement to shrink-wrapping involved modifying the smoothing step. The new smoothing operation performs edge-length-weighted Laplacian smoothing. This new technique greatly improves the output mesh by helping to prevent mesh vertices from gathering towards one particular part of the surface being parameterized. This work along with other related results was presented at the NSF/DoE Lake Tahoe "Workshop on Hierarchical Approximation and Geometrical Methods for Scientific Visualization" in October 2000.

Certain features remain difficult to obtain during shrink-wrapping because of the coarseness of the initial base mesh. We are currently investigating new techniques to ensure that shrink-wrapping can capture all of the features of the original unstructured mesh.

Finite Element Solutions to Laplace's Equation

Robert N. Rieben
University of California, Davis

Summary:

The purpose of this work is to develop a finite element program using the programming language Mathematica that could solve Laplace's equation with second-order accuracy using linear, tetrahedral finite elements. This program will serve as an implementation base for more general problems.

The program was developed from the ground up using the symbolic algorithms of Mathematica. The basis functions for the approximate solution are constructed using tetrahedral shape functions defined explicitly for each element (as opposed to being defined with respect to a reference element). The test problem is a unit sphere with Dirichlet boundary conditions. To demonstrate convergence, the solution was approximated on a series of ten successively refined meshes and compared to the exact solution. Plotting routines were developed.

Future additions to the program include the use of other three-dimensional elements (hexahedrons, prisms, etc.) and higher order finite elements. A time stepping routine will be developed for time dependent equations. This will lead to an adaptive mesh routine, which will restructure the elements of the mesh on the each time step based on the evolving solution.

A Multigrid Strategy for Accelerating Steady-State Computations of Waves Propagating with Curvature Dependent Speeds

Jonathan Rochez
University of California, Davis

Summary:

During the past year, we have been developing a code for steady state solutions to the eikonal equation in 2D and 3D using two differencing schemes for the multigrid method. Standard iterative methods show a quick reduction of the residual followed by a slow final convergence to the solution at high iterations. Such systems are ripe for the use of multigrid methods to speed up convergence.

Numerically, two different approaches have been investigated. The first approach organizes the calculation point by point on the grid using a finite-differencing scheme. The advantages of this approach include its speed in calculation time and ease of implementation, while its disadvantage is restriction to regular orthogonal grids. During this past year, code was developed to solve the eikonal equation on any orthogonal geometry with one or several detonation points. Curvature-based speeds for the wavefront propagation have been successfully implemented.

On a simple square grid geometry, solution rates of several sizes of problems with several multigrid V-cycles have been compared. In all cases, the speed was significantly faster for the multigrid method compared with calculation on a single grid. The finer the resolution, the greater the speed-up observed. For the simple geometry, the analytical solution of the problem is known to be circles emanating from each detonation point. From analyzing Fourier spectra of the error, we know it to be smooth throughout the calculation, implying the well suitedness of multigrid.

The second approach is based on discontinuous Galerkin method finite elements in a zone by zone calculation. The advantages include unstructured grids for the formation of non-regular geometries and fewer data points, while the disadvantages include larger calculation times and a less straightforward implementation. No results have been obtained yet since this second code is still being developed.

New Means of Compressing Images Using Wavelets

Joshua Senecal
University of California, Davis

Summary:

As computers continue to become more powerful the simulations run on them become larger. The result is the amount of data generated by these simulations—already on the order of terabytes—is also increasing. With the specific goal of storing the PPM dataset from the Richtmyer–Meshkov execution that won the 1999 Gordon Bell Best Performance Prize, we are developing a lossless compression method that achieves better compression than anything currently available.

Our method for compression involves transforming the data by a linear wavelet transform, and then compressing the resulting coefficients. Two observations about wavelet coefficients are these: most of the values lie at or near zero, and all the bits after the coefficient's leading bit position are random in appearance. We exploit this by compressing the position of the leading bit of the coefficient's magnitude. All zeros before this position are ignored, and all data after the leading position are copied as is. To encode the leading bit position we use Huffman encoding. Huffman's original algorithm generates codes based on the global probability of any particular symbol occurring. In our program we use context to improve compression. For example, if our program is currently processing a coefficient with a leading bit position of 1, the next coefficient is more likely to have a leading position of 0 or 1, and less likely to have a leading position of 7 or 8. We therefore construct Huffman tables based on these conditional probabilities.

As an example of how we encode, say that we need to compress the coefficient –5. In binary it is represented by a sign bit followed by an 8-bit magnitude: 100000101. Assuming that we already have a Huffman table generated, encoding would proceed as follows:

- Identify the leading bit position of the magnitude, which is 3 in this case.
- Look up the corresponding binary Huffman code for that position, say 10 in this case.
- Write out the encoded coefficient by first writing the position code (10), the sign bit (1), and everything that followed the leading bit position (01). The encoded coefficient is therefore 10 + 1 + 01 = 10101. The coefficient has been reduced to 5 bits.

As our benchmark we use the common compression program gzip. Four images, converted to 8-bit grayscale, were used in testing our compression program. One was a photograph, one was a black-and-white comic, and the other two were scientific visualization images. When compressing the photograph our program beat gzip, but gzip gets better compression on all the other images. Other tests show that our program appears to consistently do better than gzip when compressing images (like photographs) that have a lot of variation in them. gzip does better if the images have a small number of colors, contain a lot of regularity, or have large areas containing one color.

This program is still in its early stages, and more improvements are needed. Future work on improving the compressor will involve investigating additional ways of generating better Huffman codes based on context, and on improving our method so that it works better on a wider range of input

Parallel Algebraic Multigrid Methods

Marc Alexander Schweitzer
University of Bonn

Summary:

Multigrid methods are known to be fast solvers for elliptic PDEs. Specifically for strongly elliptic symmetric PDEs of second order, optimality of the convergence rate is guaranteed under weak regularity assumptions or even without regularity assumptions for the multiplicative and additive forms of the algorithm. However, for more general elliptic problems with anisotropic or discontinuous coefficient functions, standard multigrid methods using linear interpolation, full weight restriction and full coarsening lead to poor convergence rates, which are moreover dependent on the coefficient functions. Therefore, an asymptotic estimate for the condition number of the preconditioned system may be only of limited use in applications. This missing robustness has led to modifications of the standard multigrid method, based on modifications of the interpolation and restriction operators, use of multiple coarse grids, specially adapted coarse grids, or combinations of these techniques. When the interpolation operators are chosen dependent on the fine grid matrix (in a specific way) then in one dimension multigrid is just cyclic reduction, i.e., a direct solver, and in two dimensions (together with an ILU-type smoother) a robust method results. The so-called algebraic multigrid method (AMG) due to Ruge and Stüben also belongs to this class of multigrid methods, but allows for additional flexibility, through the choice of coarsened grids in an operator dependent way. Then a simple pointwise smoother is sufficient for a robust convergence behavior for a large class of problems. The basic idea of an algebraic multigrid (AMG) solver is the use of operator/matrix dependent coarsening and interpolation to achieve robustness of the solver.

The coarse grid selection within the AMG by Ruge and Stüben is inherently sequential. Numerous parallel coarse grid selection algorithms have been proposed to overcome this problem. Most approaches try to fix this problem by localizing the coarse grid selection to local process data and applying an additional setup phase to match the resulting coarse grids at process boundaries. The quality of the coarse grids generated by these approaches are dependent on the quality (with respect to AMG) of the underlying data partition of the linear system. Not only is the quality of the coarse grids affected by the data partition, but the scaling behavior of these approaches may be significantly limited by data partitions, which do not respect the physics of the underlying PDE and the discretization process.

Summary (continued):

We have proposed applying parallel multilevel graph-partitioning heuristics to the partitioned linear system to recapture the physics in the system and re-partition the data accordingly. The reduction in run time by applying AMG on an appropriate data partition though may be too small to justify this re-partitioning step by default since this pre-setup step also comes at a certain cost. For instance data partitions coming from classical domain decomposition approaches and discretizations on uniform grids and constant coefficient problems already capture the physics of the underlying problem and therefore a re-partitioning of the system will not improve the quality of the partition significantly. Modern adaptive discretization methods (adaptive FEM, meshless methods, etc.), however, usually provide the data partition either by using graph-partitioning techniques or even by some other spatial ordering technique (space-filling curves, etc.) which lead to more general data partitions.

Hence, the design of a cheap automatic mechanism to determine the 'quality' of a data-partition is necessary. The first step though towards such an algorithm is an algebraic definition of 'quality' for this purpose. To do so we have to study the effect of 'ill-partitioned' data on the quality and scalability of AMG. Here, the test problems for AMG have to be complex enough. We are working on a parallel port of a generalized (three-dimensional non-uniform discretizations) test suite to study the quality of (sequential, additive) AMG on general anisotropic diffusion and convection-diffusion problems. Furthermore, we are working on a parallel meshless method using space-filling curves for the data-partition and the interface to the hypre library.

Modeling Protein Production and Inhibition from DNA Regulation

Bahrad Sokhansanj
University of California, Davis

Summary:

This year marked the successful completion of the Human Genome Project and other whole-genome sequencing projects for rice, the fruit fly, and countless microbes. The result of these projects is the DNA sequence of almost all the genes of these organisms, and the sequences of all the protein molecules that their cells can produce. There are also small sequences of non-coding DNA which provide sites for protein binding and promote or inhibit the gene activity. However, this information is akin to a parts list for an electrical circuit: there is no wiring data, and because of limits to our current biological knowledge the function and quantity of the parts is also unknown. A major effort in bioinformatics has been to computationally predict the function of all the proteins and the wiring diagram of protein and DNA interactions. But, living cells are dynamic, constantly changing the kinds and quantities of protein molecules that they produce in response to environmental conditions: static information needs to be applied in kinetic simulation.

Our goal is to develop a dynamic simulation of bacterial gene regulation. Regulation in humans and other higher organisms uses the same fundamental processes, while adding more complex details. The model will include the binding of inhibition and activation proteins to DNA, the transcription of message RNA from the gene, and translation of the RNA to its protein product. Recent work has shown that this process depends heavily on stochastic effects resulting from the small number of molecules involved, so the traditional deterministic chemical kinetics approach is inadequate. Molecular dynamics is unsuitable for long time scales, so we are applying Monte Carlo methods with discrete numbers of molecules. Data from the literature and experiments at the Biology & Biotechnology Research Program (BBRP) are used for parameters and model validation. The long-range goal of this project is to dynamically model the genetic regulation of virulence in pathogens, and hopefully lead to a model of a whole bacterial cell.

ALE Computations Using Augmented Lagrangian Grid Speeds

Jay F. Thomas
University of California, Davis

Summary:

The Arbitrary Lagrangian–Eulerian (ALE) concept has been successfully applied to the simulation of a variety of continuum flow problems, including impact phenomena, material processing, and fluid-structure interaction. The ALE concept combines the Lagrangian approach, in which the computational mesh moves with the material, with the Eulerian approach, wherein the mesh is fixed and the material flows through it. The Lagrangian approach simplifies application of boundary conditions and tracking of material interfaces, but some flows may severely compress and distort the mesh. This causes the time step used to advance the solution to become extremely small and reduces solution accuracy in the distorted region. The Eulerian approach allows arbitrarily large deformations and turbulent flow, but material interfaces and certain types of boundary conditions may require special accommodation. The ALE concept attempts to circumvent these difficulties and obtain the advantages of both approaches by allowing the mesh to be moved in an arbitrary manner. The issue becomes one of specifying a mesh motion satisfactory for accuracy that can be efficiently generated.

Methods for specifying mesh motion can roughly be classified as static or dynamic. Static methods rely on computing the grid point locations at any time level by solving a steady grid equation (e.g., elliptic grid generation), and then interpolating the solution from the old mesh to the new mesh at a fixed time. Dynamic methods, by contrast, directly evaluate grid velocities via any appropriate law. The grid velocity equation and the physical equation may be solved simultaneously for the physical solution and the mesh. Interpolation of dependent variables from the old mesh to the new mesh is unnecessary.

ALE implementations have typically used a static approach to mesh motion as typified by mesh relaxation. Mesh relaxation is based upon elliptic mesh generation methods. The coordinates of the nodes are formulated as the solution to an elliptic problem and Jacobi iteration is used instead of a direct solution method. This results in a stencil that can be applied efficiently over the mesh. Only a few sweeps of Jacobi are used at each time step. The method thus indirectly specify the mesh node velocities but it does not lend itself to analysis in this regard.

The focus of the current effort is to investigate methods for directly specifying grid velocities in an ALE context. The general form chosen for the grid velocity consists of the fluid velocity augmented by an additive “correction” term. The idea is that the fluid velocity provides general solution adaptivity while the correction term can be chosen to control, for example, excessive mesh compression. Computations for Burgers’ equation and the 1D Euler equations, using Godunov’s method for the physical equation, have been performed using a correction term which measures how well a given weight function is equidistributed over the mesh. The approach appears effective for addressing mesh compression. To date a constant weight function has been used. Obvious extensions include use of a velocity-dependent weight function. The resulting grid velocity equation is parabolic. A significant issue for extension of the method to multiple dimensions is efficient solution of this equation, since it must be solved at each time step and stability considerations indicate that an implicit solution technique be used. The next phase of the investigation will address the extension to multiple dimensions.

ROSE Meta-program for Generating Optimizing Preprocessors

Danny Thorne
University of Kentucky

Summary:

The Overture object oriented (C++) framework for solving PDEs is hard for a compiler to optimize because it is designed with many high-level abstractions. Compilers are not capable of understanding the semantics of these high-level abstractions. ROSE is designed to be taught about the semantics of these high-level abstractions by way of specialized grammars corresponding to the code to be optimized. Thereby, ROSE will be able to generate a preprocessor that is capable of introducing optimizing transformations involving the high-level abstractions. The resulting source code, after the transformations applied by the ROSE generated preprocessor, should then be such that a compiler can build more efficient target code. One of the primary examples of the application of ROSE in Overture is to the array class library, A++/P++, from which all of Overture is derived. This library makes the develpement and use of Overture much more convenient than would be possible with primitive language constructs. However, the high-level abstractions that give rise to this convenience for the developer also have the effect of disabling the compiler from producing efficient target code. ROSE will be able to automatically transform many, if not all, uses of the array classes in a given code into something semantically equivalent but based on primitive language constructs and, hence, compiler-friendly. Since essentially all of Overture is built from the A++/P++ library, applying ROSE to transform the uses of A++/P++ classes in this way will provide for much more efficient Overture codes.

One shortcoming of ROSE, currently, is that it doesn't handle C++ templates. Many C++ codes today use templates. Hence, it is important for ROSE to be able to support templates. My main objective has been to devise and explore potential solutions to the problem of supporting templates in ROSE and to implement the support of templates if possible.

Three approaches for handling templates in ROSE have been attempted. The first two had to be discarded after much experimentation and analysis. The third is under development and promises to be successful. ROSE will be developed into a general-purpose tool for automatically introducing optimizing transformations involving the high-level abstractions in an object oriented (C++) code.

Error Control and Parallel Adaptive Grid Refinement for Convection-Diffusion-Reaction Problems

Stanamire Z. Tomov
Texas A&M University

Summary:

Simulation of flow and transport in porous media gives rise to extremely large-scale computations, making the consideration of iterative solvers natural. The solutions of such problems exhibit local nonsmooth behavior. Therefore it is essential that local grid refinement, based on a posteriori error analysis, be applied. Also, fundamental physical limitations on the computer processing speed may require the exploitation of parallelism. The goal is to create a tool that is based on discretization techniques utilizing finite elements/volumes, efficient preconditioning (in parallel) of the resulting sparse system, error control, and adaptive grid refinement.

Our computational approach is described as follows. We use mesh generator (Triangle for 2D meshes and NETGEN for 3D meshes) to generate a good coarse mesh. Then the problem is solved redundantly on the coarse mesh by every processor. The solution is used to compute a posteriori error estimates, which are used as weights in an element-based splitting of the coarse mesh into subdomains, using MeTiS. This splitting insures that the local refinements that follow will produce a computational mesh with number of triangles or tetrahedra balanced over the subdomains. Every subdomain is mapped to a processor. Then, based on a posteriori error analysis, each processor refines its region independently. After every step of independent refinement there is communication between the processors in order to make the mesh on that level globally conforming.

Concerning the a posteriori error analysis, I worked with Dr. Raytcho Lazarov on article "Error Control and Adaptive Grid Refinement for Convection-Diffusion-Reaction Problems in 3D." The article contains the description of an adaptive numerical technique based on finite volume approximations and the computational results of various model simulations of steady-state single phase flow and transport of passive chemicals in nonhomogeneous porous media in 3D.

I developed a 2D code with functionality as described in the computational approach above. The multilevel structure is used to define multigrid preconditioners. I worked closely with Panayot Vassilevski and Charles Tong on connecting this software to the Hypre Preconditioner Library. The connection is implemented using The Finite Element Interface (FEI) specification, which provides a layered abstraction that minimizes concern with internal details in the Hypre library. The initialization is done in parallel. GUI, using Motif, has been developed to utilize the selection of different Hypre options. After the solution is obtained on certain level it is sent directly through the AFJNET socket to the visualization tool SG, which resides on the user's machine. The benefit in this strategy is that the parallel machine (usually remote) is used only for computations and the local machine handles the visualization. This idea is very efficient for real time visualization of time dependent problems. The 3-D version of the code has the same features. Still under construction is the parallel local refinement step of communication between the processors for making the mesh on certain level globally conforming.

Raviart-Thomas zero-order (RTO) finite elements have been added to both the 2D and 3D codes. Under development is discontinuous approximation of convection terms for mixed finite elements. Both the 2D and 3D codes are written in C++. The parallel computations are done using the Message Passing Interface library (MPI). I finished with Raytcho Lazarov and Panayot Vassilevski article on "Interior penalty discontinuous approximations of elliptic problems," which has been submitted to *SIAM J. Scientific Computing*.

Interactive Visualization of Three-Dimensional Adaptive Mesh Refinement Data

Kevin Vlack
University of Illinois

Summary:

Adaptive Mesh Refinement (AMR) is a popular approach in scientific computing for focusing computational resources on regions of a problem where small step sizes in time and space are necessary for accurate calculation. While AMR is a natural solution for efficiently integrating stiff PDEs, it produces complicated datasets which are difficult to analyze without the aid of customized visualization packages. The main objective of the summer was to design and implement an efficient and practical volume-rendering module to interactively visualize 3D datasets created with an AMR algorithm.

There are many inherent characteristics of AMR that complicate the volume rendering process. A problem domain under an AMR framework is represented as a hierarchy of overlapping grids of increasing resolution, and each level of refinement completely contains the next, that is, any particular region of a refined grid is also represented within a coarser grid. Visualization entails selecting a mutually exclusive subset of grids from the data, and presenting the data in an interpretable manner. The situation also arises in which the scientist wishes to view overlapping areas in numerous refinement levels in order to analyze the relationship between them.

The volume renderer was implemented for X Windows in C++, using the Motif and BoxLib APIs. BoxLib is an AMR framework developed by the CFD group at Lawrence Berkeley Laboratory and is currently in use by researchers in both LBL and LLNL. The renderer uses the "splatter" technique, which approximates a voxel as an orthogonally projected polygon in order to speed the rendering process. It has been integrated into a prototype version of PAMRVIS, a parallel AMR visualization tool developed by Vince Beckner at LBL and expanded upon by Vlack in 1999.

The volume renderer was able to perform at interactive frame rates for numerous available AMR datasets, including a representation of a shear layer mixing and of a jet stream vortex. Developments were also made to the GUI of the application in order to intuitively change the parameters of the visualization in order to visualize various qualitative traits of the dataset. Although time constraints limited in-depth experimentation, basic timing techniques revealed that the maximum speedup achieved from parallel processing was about 2.8 when using four processors.

Future plans include integrating the volume-rendering module and additional interface features into the working version of the parallel visualization tool, PAMRVIS, which is used throughout the BoxLib community.

A Reduced Grid Method for a Parallel Global Ocean General Circulation Model

Michael Everett Wickett
University of California, Davis

Summary:

A limitation of many explicit finite-difference global climate models is the timestep restriction caused by the decrease in cell size associated with the convergence of meridians near the poles. A computational grid in which the number of cells in the longitudinal direction is reduced toward high-latitudes, keeping the longitudinal width of the resulting cells as uniform as possible and increasing the allowable timestep, is applied to a three-dimensional primitive equation ocean-climate model. This "reduced" grid consists of subgrids that interact at interfaces along their northern and southern boundaries, where the resolution changes by a factor of three. Algorithms are developed to extend the finite difference techniques to this interface, focusing on the conservation required to perform long time integrations, while preserving the staggered spatial arrangement of variables and the numerics used on subgrids. The reduced grid eliminates the common alternative of filtering high-frequency modes from the solution at high-latitudes to allow a larger timestep and reduce execution time per model step by roughly 20 percent.

Our reduced grid model is implemented for parallel computer architectures with two-dimensional domain decomposition and message passing, with speedup results comparable to those of the original model. Both idealized and realistic model runs are presented to show the effect of the interface numerics on the model solution. First, a rectangular, mid-latitude, flat-bottomed basin with vertical walls at the boundaries is driven only by surface wind stress to compare three resolutions of the standard grid to reduced grid cases, which use various interface conditions. Next, a similar basin with wind stress, heat, and fresh water forcing is used to compare the results of a reduced grid with those of a standard grid result while exercising the full set of model equations. Finally, global model runs, with topography, forcing, and physical parameters similar to those used for ocean-climate studies, are advanced to a near equilibrium state for both the reduced grid and the standard grid. Differences between the two are presented for typical fields of interest, and very little degradation of the solution due to the reduced grid is observed.

Integrating Arches with PETSc

Wing K. Yee
University of Utah

Summary:

Arches is a Computational Fluid Dynamics code being developed by Dr. Phil Smith and his colleagues at the University of Utah. The motivation for the development of this code is to simulate large-scale fire. First-principles simulations of fire are extremely complicated, taking into account mixing, reaction, and radiation, in addition to complex fluid mechanics. For large-scale fire, not only are the variables resolved on the grid scale important, the subgrid level variables are also important and need to be modeled. The resulting multiple-scale nonlinear partial differential equations need to be solved accurately on each time step. DOE's PETSc (the Portable, Extensible Toolkit for Scientific computation) is a software environment with solvers for discretized PDEs in parallel or serial. It has been employed to port legacy CFD codes to the ASCI parallel platforms, and one such port resulted in a 1999 Gordon Bell Prize. Flexibility is a trademark of PETSc. It has a Newton method base solver, but offers many different options for preconditioning and matrix construction. It is also compatible with all three major high-level languages: FORTRAN, C, and C++. The price of working with PETSc is allowing it to control the major distributed data structures, and supplying routines that link the physics-laden user code in local (legacy) data structures with the PETSc data structures.

Our first step to integrate PETSc into Arches involved an older version of Arches, Banff. This has provided valuable guidance for the development of Arches. From this learning experience, a programming standard for the Arches code has been set. The tests with Banff were done on the ASCI Blue machine.

In addition to the parallelization project, a new Monte Carlo numerical integration method for the mixing model was also explored.

Scalable Adaptive Multilevel Algorithms

Gerhard Zumbusch
University of Bonn

Summary:

The solution of boundary value problems of partial differential equations is considered. The goal is to construct an efficient solver, which uses an optimal order solution algorithm (multigrid), a low number of unknowns (adaptivity), very little memory (hash storage), and runs on large parallel computers. It is mainly the grid adaptivity that poses some difficulties for the parallel implementation.

Multigrid and multilevel methods are optimal order solution algorithms for equation systems stemming from the discretization of PDEs. They require linear time, that is O(n) operations for n unknowns. In the presence of singularities of the solution the convergence of standard discretizations is not as rapid as for smooth, regular solutions. One way to overcome this problem is adaptivity. The grid is refined during the computation, where indicated by an error estimator or error indicator.

The components of our adaptive multigrid code have to be parallelized. A single data decomposition is used on the distributed memory computer. However, each component requires a different treatment. A dynamic load-balancing scheme is used for the parallel execution of the code. It is based on space-filling curves. The parallelization technique has also been applied to sparse grid discretizations.

ISCR

Institute for Scientific Computing Research

Workshop and Conference Reports

CASC Center for Applied Scientific Computing

May 2000

Synopsis of Workshop Events:

Power Programming Short Course

The Institute for Terascale Simulation, an ASCI-supported arm of the ISCR, organized a 3-day *Power Programming Short Course* to enable laboratory code developers to come to grips with performance-oriented aspects of parallel computation on distributed shared memory machines, such as the ASCI White machine that arrived at the laboratory in July 2000. The instructors were Steve White (IBM), Larry Carter (UCSD), David Culler (UCB), ISCR collaborator Clint Whalley (University of Tennessee), and Bill Gropp (Argonne National Laboratory). Sixty-five people (LLNL training center room capacity) attended—most for the full three days, May 15-17, 2000.

The topics were as follows:

- POWER3 Architecture and Tuning for the ASCI White system: the POWER3 microarchitecture, tuning advice and experience, and ASCI White overview (White)
- Performance Programming, I: exploiting the Power processor in the design of kernels and data structures for scientific applications (Carter)
- Performance Programming, II: cache and TLB issues (Carter) Sensitivity-based Performance Analysis Tools: understanding performance thresholds through parameterized microkernels (Culler)
- High Performance Communication, I: MPI-1, point-to-point and collective (Gropp)
- High Performance Communication, II: parallel I/O, MPI/OpenMP tradeoffs, and communication monitoring tools (Gropp)

The workshop was of value beyond the ASCI context, since the Power architecture is in many ways typical of other hybrid architecture high-end scientific computers. Participants were encouraged to bring laptops or overhead slides with which to display code or precisely state questions or comments.

The short course was organized by David Keyes of the ISCR, John May of CASC, and Mary Zosel of Livermore Computing.

For more information, see the website
http://www.llnl.gov/CASC/workshops/course_reg/short_course_info.html

Synopsis of Workshop Events:

The Second Workshop on Mining Scientific Datasets

The Second Workshop on Mining Scientific Datasets was held July 20-21, 2000, at the Army High Performance Computing Research Center (AHPCRC) in Minneapolis, Minnesota. The goal of the workshop was to bring together researchers from the data mining and scientific computing communities in order to better understand how data mining can be used for the exploration of scientific datasets. Whereas the literature of data mining is dominated by demographic, economic, and related applications, the participants concentrated on the challenges and opportunities that are unique to the application of data mining to scientific data sets.

This workshop followed an earlier workshop on the same topic that was held in September 1999 at AHPCRC. The first workshop covered a broad range of topics with the aim of identifying the main disciplines that could help address the challenges in mining scientific data sets. The second workshop focused on how the research being done in these diverse technical areas could be effectively harnessed to solve the problems of scientific data analysis.

The 22 talks at the workshop were broadly divided into two categories:

- Scientific applications where data mining techniques were either being applied, or could be applied. These included areas such as astronomy, physics, earth sciences, fluid dynamics, protein folding, and spatial data sets arising in Geographic Information Systems.
- Data mining algorithms, including traditional pattern recognition techniques such as clustering, classification, and association rules, as well as various data pre-processing techniques such as dimension reduction, feature extraction, and feature selection.

A few speakers gave a broad overview of the field and shared their experiences with mining scientific data. Jagdish Chandra, Professor at George Washington University and former director of the Math and Computer Science Division at the Army Research Office, described the mathematical challenges in understanding high-dimensional data sets. He identified several common features of current data sets, including their massive size (terabytes and petabytes), their complex structure in terms of the relations between different parts of the data, their noisiness as a result of the way in which the data was collected, as well as our imperfect understanding of the basic processes for which we seek information.

Chandra observed that questions regarding the structure of the data had been present for centuries. However, it was the scale of the problem, as well as its internal complexity, that required us to rethink old ideas and solve the problem with an infusion of new ideas. He characterized these “model rich” data mining problems by their non-linearity as well as intrinsic uncertainty. He believed that computational power by itself was not enough to solve these problems—a deep understanding of the natural and mathematical structure underlying the new data was necessary in order to extract information with scientific, technological, and societal implications. He echoed the views of several experienced data miners when he observed that, to address these problems, we need a multi-disciplinary approach, combining the expertise of mathematicians, numerical analysts, signal processing engineers, statisticians, and computer scientists.

Synopsis of Workshop Events (continued):

A different perspective was given by Padhraic Smyth from University of California at Irvine, who described some of the new research ideas being pursued in data mining. He observed that data mining could be considered as the data driven discovery of models and patterns from massive observational data sets. This would require a modern-day data miner to be proficient in statistics, modeling, optimization, and search, as well as algorithm design and data management. The process of data mining could be considered as being composed of several components, including the dataset itself, the task (e.g., prediction, clustering, etc.), the representation of the problem (i.e., the underlying language used), the score function and the optimization technique used, and the data access approach. Each of these components would be driven by the application. Smyth identified three areas of new research currently being pursued in data mining—scalability either by scaling down the data or scaling up the algorithms, the identification of patterns vs. models, and the clustering of objects of different dimensions.

A broad range of application areas was covered by the workshop participants. The problems being tackled by means of data mining techniques included the detection of coherent structures in turbulent boundary layers, the prediction of three-dimensional contact potentials among protein residues, the identification of galaxies with a bent-double morphology, the content-based querying of earth science data, as well as the Virtual Observatory concept being proposed by astronomers as they mine several large-area sky surveys. In the area of algorithms, the topics included the generation of long patterns, dimension reduction of image feature descriptors, clustering via the construction of a decision tree, as well as techniques and protocols for mining distributed data.

The workshop included a panel discussion on spatial data mining. The panelists, led by Sashi Shekhar (University of Minnesota), included Jiawei Han (Simon Fraser University), James LeSage (University of Toledo), Suchi Gopal (Boston University), and Sanjay Chawla (Vignette Corporation). The discussion centered around the challenges encountered in mining spatial data, including the appropriate ways to discretize a spatial domain, the selection of features in light of the diversity of possible spatial relationships, the metrics for evaluation that capture spatial accurately, etc.

During the workshop, several talks touched upon the problems often encountered in mining scientific datasets, including:

- the heterogeneity of the data
- the high dimensionality of the feature space
- the problems in extracting relevant features from simulation data
- the spatio-temporal nature of the data
- feature selection
- scaling the algorithms to massive data sets
- the presence of noise in the data
- the lack of ground truth

Synopsis of Workshop Events (continued):

Dr. N. Radhakrishnan, the Director of the Corporate Information and Computing Directorate of the Army Research Laboratory, closed the workshop with encouraging words for the participants, urging them to address the challenges that remain in applying data mining techniques to massive scientific data sets resulting from computer simulations and observations.

The overwhelming interest expressed by the 110 attendees of this second workshop has led to a follow-on workshop that will be held in conjunction with the first SIAM International Conference on Data Mining. Additional details on this conference are available at *http://www.siam.org/meetings/sdm01/*.

The workshop was sponsored by AHPCRC, the Minnesota Supercomputing Institute, and the Center for Applied Scientific Computing at the Lawrence Livermore National Laboratory. The co-organizers were Bob Grossman (University of Chicago and Magnify, Inc.), Chandrika Kamath (Lawrence Livermore National Laboratory), Vipin Kumar (University of Minnesota), and Raju Namburu (Army Research Laboratory).

A website with details on the first and second workshops, together with abstracts and presentation slides, is at *http://www.ahpcrc.org/conferences/*

Contributed by Chandrika Kamath, Center for Applied Scientific Computing, LLNL.

Synopsis of Workshop Events:

Workshop on Solution Methods for Large-Scale Nonlinear Problems

The solution of large-scale, fully coupled multiphysics models is vital to progress in computational science. Efficiently computing solutions to the mathematical equations underlying these models requires effective numerical algorithms for solving very large systems of coupled nonlinear equations. Scalable algorithms that can exploit the power of massively parallel computers are especially necessary. Several successful approaches to constructing such algorithms have emerged in recent years, but the need to treat more complex models and problematic physics continues to drive research.

The Workshop on Solution Methods for Large-Scale Nonlinear Problems, held July 26–28, 2000, at the Four Points Hotel in Pleasanton, CA, brought together many of the most active researchers in nonlinear solution algorithms and applications. The workshop was hosted by the Center for Applied Scientific Computing and the Institute for Terascale Simulation at Lawrence Livermore National Laboratory (LLNL). Major themes included Newton–Krylov methods, nonlinear multigrid methods, preconditioning techniques, operator-split and fully implicit schemes, and large-scale sensitivity analysis and optimization. Major application areas included general fluid dynamics, combustion and other reacting flows, porous media flows, radiation diffusion problems, and design optimization.

The workshop was a distant follow-on to similar workshops on large-scale nonlinear problems held at Utah State University in 1989 and 1995. In the spirit of those earlier workshops, the schedule of talks and discussion sessions allowed liberal break time to encourage informal interactions among the participants. Workshop attendees came from academia (25), government laboratories (21), and industry (1) and were mainly from the United States.

Each of the three workshop days emphasized particular subject areas, with nine presented talks followed by a moderated discussion of selected issues. The first day focused on Newton–Krylov methods and nonlinear multigrid, also called the Full Approximation Scheme (FAS); the second day focused on applications and problem formulation; and the third day focused on algorithmic issues, sensitivity analysis, and optimization.

The first day began with two talks on techniques for enhancing the convergence of Newton–Krylov methods. The first introduced a novel nonlinear preconditioning method based on Schwarz domain decomposition; the second outlined approaches based on the singular value decomposition. Later presentations addressed applications of Newton–Krylov and nonlinear multigrid methods to combustion, groundwater flow, and radiation diffusion problems. The final talk addressed the benefits of using automatic differentiation instead of finite differences in computing matrix-vector products in Newton–Krylov methods. To conclude the day, Van Henson and Jim Jones led a moderated discussion on the differences in performance between Newton's method and FAS on various problems. The basic iteration scheme of FAS was discussed along with the full multigrid version. Methods using multigrid to solve Jacobian systems for each Newton iteration (Newton-multigrid) were presented. The consensus of attendees who had experience with FAS indicated that FAS requires less memory and has a larger basin of attraction than Newton-multigrid methods, but that it is harder to get FAS to work on new problems.

Synopsis of Workshop Events (continued):

The second day, which emphasized applications and problem formulation, included presentations on nonlinear solvers for reacting flows, multiphase groundwater flow, chemical reactions, radiation diffusion problems, powder consolidation, and Einstein's equations of general relativity. Solution approaches included operator splitting techniques and fully implicit formulations, domain decomposition, and pseudo-transient continuation. Peter Brown and Carol Woodward led the moderated discussion that afternoon on whether or not to operator-split coupled systems of nonlinear equations. Participants generally agreed that the holy grail is a fully implicit solve of all problem physics, but simulation technology has not evolved enough to handle all the necessary complexity. Significant progress has been made, however, in development of solvers and methods for these fully implicit formulations.

The third day, which focused on algorithmic issues, sensitivity analysis, and optimization, included presentations on nonlinear elimination methods, performance "stress points" of parallel implicit solvers, techniques for finding solution sensitivities to problem parameters, and formulation and solution techniques for nonlinear PDE-constrained optimization. Applications included aerodynamic analysis and design optimization as well as optimal control of Navier–Stokes flows. Steven Lee and David Keyes moderated the afternoon discussion on the role of sensitivity analysis and optimization in scientific computing. Discussion participants pointed out that aerodynamic and automotive design were fields where sensitivity analysis and optimization were now being used routinely and that, with the development of effective automatic differentiation tools, more fields will begin using these techniques in the future.

From the discussion and presentations, it became clear that the field of solution methods for large-scale nonlinear problems has evolved considerably in the five years since the last workshop. Although Newton–Krylov and Newton–Krylov–Schwarz methods were important topics at the 1995 meeting, they were much more prominent in this workshop. This change reflects advancement of the Newton–Krylov family of methods and the current widespread acceptance of these techniques as the methods of choice in many difficult applications. Similarly, fully implicit and fully coupled problem formulations were much more in evidence in this workshop as a result of advances in solution technology that have made fully coupled formulations feasible in previously intractable applications. Such advances have also made it possible to more effectively treat difficult PDE-related optimization problems, such as design optimization. Accordingly, the optimization talks in this workshop focused more on PDE-related applications and less on general optimization algorithms.

A notable new feature of this workshop was the work on nonlinear multigrid, which did not appear at all in the previous workshop. Although the

Synopsis of Workshop Events (continued):

method is not new, it has not been widely used until recently, when advances in the solution of heterogeneous problems made the method more applicable to problems of interest. Another area not previously featured was sensitivity analysis. Even though this area has been a subject of importance for some time, the introduction of special Newton–Krylov techniques in recent years has resulted in much more effective methods for large-scale problems.

Many of the workshop topics seem likely to lead to interesting future research. Nonlinear preconditioning, which generated much discussion among the participants, shows considerable potential for resolving well-known stagnation problems associated with Newton's method. Nonlinear multigrid, though not as new, also appears to have much unexploited potential, not only as a nonlinear solver but perhaps as a nonlinear preconditioner as well. Similarly, sensitivity analysis and design optimization seem likely candidates for further attention as algorithms are improved and more challenging applications addressed. Another area, the use of operator splitting techniques for preconditioning fully implicit formulations, should lead to more accurate simulations using fully coupled approaches. This technique will also allow the reuse of current methods and codes to solve these preconditioning systems. Nonlinear elimination, which was observed to be useful in treating shocks in compressible flow applications, also appears worthy of further attention. Finally, automatic differentiation seems at last to have emerged as a practical tool that should see many applications in Newton–Krylov methods. Automatic differentiation has the potential to alleviate algorithmic slowdowns and failures that sometimes occur when finite-difference techniques are used in approximating matrix-vector products.

A special issue of the journal, *Numerical Linear Algebra with Applications*, will feature papers from this workshop. The issue will be co-edited by workshop organizing committee chair Carol Woodward and journal editor Panayot Vassilevski. On the basis of the enthusiasm of workshop participants, the apparent algorithmic advances, and the challenge of increasing problem scale, LLNL's Center for Applied Scientific Computing and Institute for Terascale Simulation will sponsor more workshops in this series in coming years.

For more information, see the website
http://www.llnl.gov/CASC/workshops/workshop_reg/workshop072600_info.html

Contributed by Homer Walker, Mathematical Sciences Department, Worcester Polytechnic Institute and Carol S. Woodward, Center for Applied Scientific Computing, LLNL.

Synopsis of Workshop Events:

Computational Science Graduate Fellows Conference

The Krell Institute, which operates the *Computational Science Graduate Fellowship* (CSGF) program for the Department of Energy, organizes a summer research symposium for fellows and their advisors from leading universities across the United States. In 2000, for the first time, Krell decided to host the symposium on site at the three DOE research laboratories in the San Francisco Bay area. The ISCR assisted Krell with the local logistics of the three-day conference, which included on-site presentations at the Lawrence Livermore, Sandia Livermore, and Lawrence Berkeley laboratories. Speakers were provided by all three of the local laboratories, and other DOE organizations.

Seven CSGF students were interning at the ISCR already at the time of the July 27–29 meeting. They were joined by thirty-nine of their peers and a contingent of ten advisors. Thirty-six DOE scientists from all over the national Office of Science and Defense Programs laboratory complex, and DOE administrative staff from Washington, D.C., converged at the conference to provide technical content and practical career advice to the "cream of the crop" graduate students.

David L. Brown of CASC gave one of the plenary presentations at the Symposium, entitled *Overture: An Object-oriented Framework for Solving Partial Differential Equations in Complex Geometry.*

Evi Dube of B Division described *Computational Research Challenges for an ASCI Simulation Code* in one of the breakout sessions.

Another major LLNL code project, *Using AMR in CFD – from Shock Tubes to Lasers*, was presented by A Division's Jeff Greenough.

Richard Hornung of CASC presented a faculty collaborative ISCR project entitled *A Hybrid Model for Gas Dynamics that Couples Continuum and Direct Simulation Monte Carlo Methods Using Adaptive Mesh Refinement* at another LLNL breakout.

Completing the LLNL presentations was a student collaborative ISCR project, *The ROSE Project: The Optimization of Object-oriented Scientific Applications*, presented by Dan Quinlan of CASC.

For more information about the CSGF program, see the website at *http://www.krellinst.org/CSGF/*

The conference website is *http://www.krellinst.org/CSGF/conference.html*

Synopsis of Workshop Events:

Internships in Terascale Simulation Technology

In collaboration with LLNL's Science and Technology Education Program (STEP), and with sponsorship from the DOE Defense Programs Office in Washington through Beverly Berger, the ISCR organized and conducted a series of tutorial lectures to enrich the technical and social experiences of the large population of summer students interning at the laboratory. Dubbed *Internships in Terascale Simulation Technology* (ITST), this ten-week series with two lectures per week featured ten different presenters.

Led and anchored by Alice Koniges, fresh from the success of her new Morgan-Kaufman book *Industrial-strength Parallel Computing*, with six lectures, the tutorial also featured two sessions each from two other recent LLNL authors: John May, whose tome on *Parallel I/O* (also from Morgan-Kaufman) appeared in September, and Van Emden Henson, who co-authored an update of the 1987 SIAM classic, *A Multigrid Tutorial*.

Five LLNL computational physicists shared from their experiences with real-world applications related to the ASCI mission. David Brown presented LLNL's *Overture* computational environment for the solution of PDEs in complex geometry, featuring several examples from combustion, immiscible fluids, and other fields. Garry Rodrigue gave a two-lecture introduction to the numerical analysis of shock physics. Howard Scott, Alek Shestakov, and Lin Yang combined for two lectures on parallel programming techniques and data communication paradigms in practical physics codes. David Keyes, who together with Koniges organized the overall program, finished up with two lectures on high-performance parallel algorithms for PDE simulation, which included an "anatomy" of the Bell-prize-winning computations he performed with his graduate student Dinesh Kaushik and collaborators from DOE's Argonne National Laboratory and NASA.

Recently a student himself, CASC's Gary Kumfert gave a workshop on the practical matter of preparing scientific presentations. This topic was presented at the summer's midpointæafter the interns were sufficiently into their projects to think concretely in their terms but with plenty of time left before STEP's Poster Presentation day in August, at which several ISCR students showed off the fruit of their labors.

It was gratifying to find permanent CASC researchers attending an occasional subseries. For example, several CASC computer scientists attended Garry Rodrigue's lectures on numerical simulation of shocks after finding, being employed at the laboratory for a season, that they needed to understand the associated vocabulary and challenges to converse with laboratory "clients." Similarly, veteran computational researchers took advantage of John May's tutorial lectures on parallel debugging and parallel I/O.

Synopsis of Workshop Events (continued):

Schedule of Lectures

June 13

Parallel Computing Resources/Parallel Architecture Overview
Speaker: Alice Koniges

June 15

Performance Issues, Measuring and Reporting Performance
Speaker: Alice Koniges

June 20

Parallel Programming Models and Languages I
Speaker: Alice Koniges

June 22

Programming Models and Languages II, Parallel I/O, Parallel file systems
Speaker: Alice Koniges

June 27

Performance Optimization, Optimization Issues
Speaker: Alice Koniges

June 29

Case Studies: How Much Can Performance be increased in a Real Application?
Speaker: Alice Koniges

July 6

Power Presentations
Speaker: Gary Kumfert

July 11

Basic Parallelization,
Speaker: Alek Shestakov
and
Parallel Applications in Physics,
Speaker: Lin Yang

July 13

Mixed Models (Open MP or Pthreads with MPI),
Speaker: Howard Scott

Synopsis of Workshop Events (continued):

July 18

PDEs in Complex Geometry
Speaker: David Brown

July 20

PDEs in Complex Geometry
Speaker: David Brown

July 25

Numerical Shock Simulation
Speaker: Garry Rodrigue

July 27

Numerical Shock Simulation
Speaker: Garry Rodrigue

August 1

Parallel I/O and Parallel Debugging
Speaker: John May

August 3

Parallel I/O and Parallel Debugging
Speaker: John May

August 8

A Multigrid Tutorial
Speaker: Van Henson

August 10

A Multigrid Tutorial
Speaker: Van Henson

August 15

Parallel Solver Infrastructure
Speaker: David Keyes

August 17

Parallel Solver Infrastructure
Speaker: David Keyes

Synopsis of Workshop Events:

Fifth Symposium on Overset Grids & Solution Technology

Overset grid technology has emerged as an effective method for providing discretizations of complex geometries for physical simulations. Instead of creating a single grid to represent a geometry and the domain of interest, a collection of overlapping, or overset, component grids is provided; each grid represents a portion of the domain. Solving a system of equations over the entire problem involves determining a solution on each component grid and communicating the solution between grids in overlapping regions.

As the latest in a ten-year series of semiannual symposia, the Fifth Symposium on Overset Grids & Solution Technology provided an effective forum for ideas and applications. The symposium was hosted by the University of California at Davis and held at that campus on September 18th–20th, 2000. A diverse range of institutions further supported the meeting representing government, academic institutions, and industry. Sponsors included the Institute for Scientific Computing Research (ISCR) at LLNL, the U.S. Army Research Office, UC Davis Department of Mechanical & Aeronautical Engineering, and Intelligent Light Corporation, the developers of a popular engineering visualization environment.

Presentations and discussions during the symposium involved a variety of overset grid topics spanning numerical methods, grid generation, applications, and software development. The meeting was organized to provide a mixture of these topics each day, concluding with a panel discussion entitled “Pros and Cons of Overset Grid Solution Technology.” Participants represented government, academic, and industrial researchers from around the world, applying overset grids to many different applications. A diverse set of application domains were presented including interface dynamics, moving body simulations, chemical vapor deposition, tidal flow simulation, and unsteady insect flight dynamics.

Discussions regarding numerical methods were as diverse as the applications they were intended to model. Brian Miller and Bill Henshaw, members from the Overture project at Lawrence Livermore’s Center for Applied Scientific Computing (CASC), presented their development of level set methods on overlapping grids. Dan Quinlan and Bobby Philip, also from the Overture project, discussed their application of overset grids in a hierarchical manner to build adaptive mesh elliptic equation solvers. ISCR summer student Lars Carlson from Chalmers University in Sweden related his recent work on line-implicit methods for the incompressible Navier-Stokes equations, a topic made more difficult by the unstructured nature of overlapping grid connectivity. Robert Trammel, of CFD Research Corporation, presented papers on time integration techniques for unsteady flow applications and overlapping additive Schwarz methods for the Helmholtz equations. On the second day, Anders Petersson, an Overture team member, provided an investigation of pressure boundary conditions for the incompressible Navier–Stokes equations.

Grid generation remains an area of considerable activity within the overset grid community. Many of the mesh generation presentations focused on running larger, more complicated problems and devising methods to automate the generation of the collection of meshes. William Chan from NASA Ames presented two papers highlighting recent developments of his interactive mesh generator OVERGRID and his automated surface mesh generation technologies. There were two talks describing algorithms for the automated assembly of the overlapping meshes, including the creation of the necessary interpolation stencils. Representatives from industry demonstrated the capabilities of their products.

Synopsis of Workshop Events (continued):

Meng-Sing Liou, a researcher at NASA Glenn, described an interesting alternative to overlapping component meshes using a new technology that replaces the overlapping regions in an overset grid with unstructured patches. An unusual and impressive demonstration of the overlapping mesh concept was provided by Professor Nakahashi of Tohoku University, Japan, who used overset unstructured meshes for large moving body problems.

Many applications of overset grid technologies were represented at the symposium. One group utilized overlapping grids to investigate three-dimensional unsteady flow of blood in major arteries. Another researcher used the Overture software framework to simulate chemical vapor deposition. The first day saw Petri Fast of the Overture project present a method for modeling interface dynamics using overset grids. Aerodynamic applications included wind power generation, modeling the interactions of bodies in buoyant jets and plumes, and turbomachinery simulations for the space shuttle. One notable example was provided by Dora Yen, a UC Davis graduate student, modeling the use of micro-electro-mechanical (MEMs) devices to control flow over wing sections. Her work coupled the use of wind tunnel experimentation and computer simulation using overset grids to investigate this new method for controlling the flow around aircraft wings. Overset grids are often used to simulate bodies moving relative to one another within a flow field. Several examples of this class of problem were presented at the symposium, including simulations of store separation from aircraft, space shuttle booster launch modeling and the deployment of multi-warhead projectiles.

Several presentations described improved software techniques making the overset grid method more efficient and accessible to other researchers. A topic permeating many talks was the exploitation of parallel computer architectures within existing overlapping grid tools. Brian Gunney presented Overture's recent developments in this area while Cetin Kiris described NASA Ames' recent efforts. Software framework projects were also important contributors to the symposium. The Overture Project from CASC produced several talks, both from within the group itself as well as from users. Andrew Wissink, another CASC representative, presented related technology in the SAMRAI adaptive mesh framework. Variations and applications of NASA's OVERFLOW code were also prominent. Mississippi State University researchers presented another software toolkit for overlapping mesh problems. As a proponent of the technology, the Army Research Labs proposed to include overset grid software as a part of their Major Shared Resource Center.

The symposium concluded with panel discussion on "The Pros and Cons of Overset Grid Solution Technology." As evidenced by the many and varied applications presented, most panelists were already in agreement that the method was flexible and effective. Current research issues such as automated mesh generation and parallel computing were emphasized as important for the continued development of the technology. Several expressed the desire to see the techniques used in an even broader community of researchers and engineers, encouraging the development of freely available software to accomplish this end. Concluding the remarks as well as the symposium, participants received an invitation to the Sixth Overset Grid Technology Symposium, to be held in the Washington, DC, area in 2002.

For more information, see the website *http://ntserver.itd.ucdavis.edu/Chimera2000/*

Contributed by Kyle Chand, CASC

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551